
TRANSCRANIAL ULTRASOUND SIMULATION WITH UNCERTAINTY ESTIMATION

Antonio Stanzola¹, José A. Pineda-Pardo², Bradley Treeby^{1*}

¹University College London, Gower Street, London, WC1E 6BT, UK

²HM CINAC (Centro Integral de Neurociencias Abarca Campal)
Fundacio n Hospitales de Madrid, HM Hospitales, Madrid, Spain

ABSTRACT

Transcranial ultrasound simulations are increasingly used to predict *in situ* exposure parameters for ultrasound therapies in the brain. However, there can be considerable uncertainty in estimating the acoustic medium properties of the skull and brain from computed tomography (CT) images. Here, we show how the resulting uncertainty in the simulated acoustic field can be predicted in a computationally efficient way using linear uncertainty propagation. Results for a representative transcranial simulation using a focused bowl transducer at 500 kHz show good agreement with unbiased uncertainty estimates obtained using Monte Carlo.

1 Introduction

Transcranial ultrasound simulations are increasingly used to predict *in situ* exposure parameters for ultrasound therapies in the brain, including for blood-brain barrier opening [1], high-intensity focused ultrasound ablation [2], and transcranial ultrasound stimulation (TUS) [3]. Simulations are particularly important for treatment planning in TUS, as online monitoring methods (such as thermometry and acoustic radiation force imaging) are not yet sufficiently sensitive for use with low-intensity ultrasound therapies [4].

The current approach for transcranial ultrasound simulation is to map the geometry and acoustic properties of the skull and brain from computed-tomography (CT) images [5]. However, there can be considerable uncertainty in the mapping from Hounsfield units to acoustic parameters [6, 7]. Errors in the sound speed of the skull in particular, can result in changes to the amplitude and shape of the acoustic field inside the brain [8, 9, 10]. In principle, these variations could be quantified using a Monte Carlo approach, where many simulations are performed using different acoustic property mappings, and the results of these simulations used to quantify the uncertainty. However, Monte Carlo estimation can take a large number of samples (i.e., simulations) to converge, significantly increasing the computational cost of model-based treatment planning.

In general, quantifying uncertainty is a challenging task for which many methods have been developed to avoid expensive Monte Carlo simulations. These range from the application of the Koopman operator [11], to density estimation [12]. Here, we demonstrate how computationally efficient uncertainty estimation can be achieved using linear uncertainty propagation with a differentiable wave simulator, j-Wave [13].

2 Linear uncertainty propagation

For a given scanner and image acquisition and reconstruction settings, the mapping from a CT image in Hounsfield units to an image of mass density in $\text{kg}\cdot\text{m}^{-3}$ can be performed with the aid of a calibration image of a test object with known density. Figure 1(a) shows an example of such a calibration acquired using a CIRS electron density phantom (Model 062M) using the scanner details described in [14]. The error bars show the standard deviation in the image values averaged over each tissue equivalent electron density plug. A linear curve is then typically used to map from density to sound speed (or equivalently, two different linear mappings are used to map to density and sound speed from Hounsfield units [5]). Figure 1(b) shows a scatter plot of the density and sound speed of the skull samples measured

*bradley.treeby@ucl.ac.uk

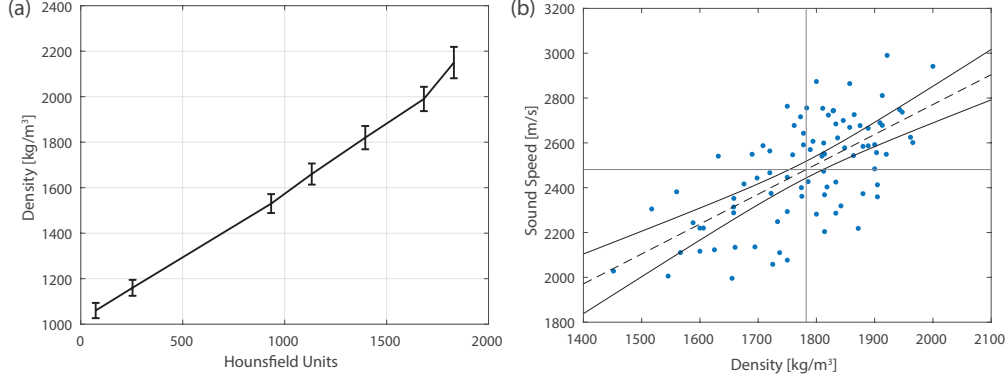


Figure 1: (a) Conversion between Hounsfield units and mass density obtained using an image of a CIRS electron density phantom (Model 062M). The error bars show the standard deviation in the image values averaged over each tissue equivalent electron density plug. (b) Conversion between density and sound speed using the data from [6, 7]. Each data point shows the density and sound speed measured for a specific skull sample (one outlier was removed from the data). The dashed line shows a linear fit to the data, and the solid lines show the 95% confidence interval for the linear model. The cross hairs show the mean values.

in [6, 7], along with a linear fit and the 95% confidence interval for the linear model. There is considerable spread in the measurement data, which may be attributed to both measurement noise (e.g., in the time-of-flight picking used to estimate the sound speed) and inherent variations in the mapping for different skulls and different regions of the skull. Here, we only consider the uncertainty in the mapping from density to sound speed, but in principle, it is also possible to propagate the uncertainty in the mapping from Hounsfield units to density, and to account for uncertainty in other acoustic and thermal parameters.

Using the data shown in Fig. 1(b), it is assumed the mapping between density ρ and sound speed c is given by the linear relationship

$$c(\alpha, \beta) = \hat{c} + \alpha(\rho - \hat{\rho}) + \beta, \quad (1)$$

where $\hat{c} = 2481 \text{ m.s}^{-1}$ and $\hat{\rho} = 1782 \text{ kg.m}^{-3}$ are the mean values for the sound speed and the density in the data, respectively (shown by the cross-hair in Fig. 1(b)). Fitting the model on mean-zero data results in a linear fit with uncorrelated estimates for the α and β parameters. While correlation of the input variables can be accounted for by linear uncertainty propagation, this simplifies the implementation. The least squares estimate of the slope and intercept are given by $\alpha_0 = 1.333$ and $\beta_0 = 0 \text{ m.s}^{-1}$, with a standard deviation of $\sigma_\alpha = 0.168$ and $\sigma_\beta = 18.8 \text{ m.s}^{-1}$.

Rather than performing a single acoustic simulation with a fixed mapping between density and sound speed (i.e., just using the values of α_0 and β_0), to account for the uncertainty in the mapping, it is assumed that α and β can be represented by uncorrelated Gaussian probability distributions. For a fixed problem (for example, a particular transducer and density map), the simulated acoustic pressure field p can be written as a function of the parameter vector $\mathbf{v} = (\alpha, \beta)$

$$p = f(\mathbf{v}) . \quad (2)$$

Here, f includes the wave simulation, and any other pre or post processing steps needed to generate the output pressure field. To propagate uncertainty, it is assumed that \mathbf{v} is drawn from a normal distribution centered around μ_v with variance Σ_v , given by

$$\mu_v = (\alpha_0, \beta_0), \quad \Sigma_v = \begin{pmatrix} \sigma_\alpha^2 & 0 \\ 0 & \sigma_\beta^2 \end{pmatrix}, \quad (3)$$

that is, $\mathbf{v} = \mu_v + \varepsilon$, with $\varepsilon \sim \mathcal{N}(0, \Sigma_v)$. If ε is small enough, Eq. (2) can be expanded as a Taylor series up to second order

$$f(\mathbf{v}) \approx f(\mu_v) + J(\mathbf{v} - \mu_v), \quad \text{with } J = \nabla f|_{\mathbf{v}=\mu_v} = \begin{pmatrix} \partial f / \partial \alpha \\ \partial f / \partial \beta \end{pmatrix}, \quad (4)$$

where J is the Jacobian.

The truncated first-order Taylor series expansion of $f(\mathbf{v})$ is a linear function of \mathbf{v} , thus inputs drawn from a normal distribution are mapped to outputs that are also normally distributed. The mean-value of the distribution is $f(\mu_v)$ (in other words, the simulated pressure field using α_0 and β_0), while the variance is given by [15]

$$\text{Var}(f(\mathbf{v})) = J \Sigma_v J^T . \quad (5)$$

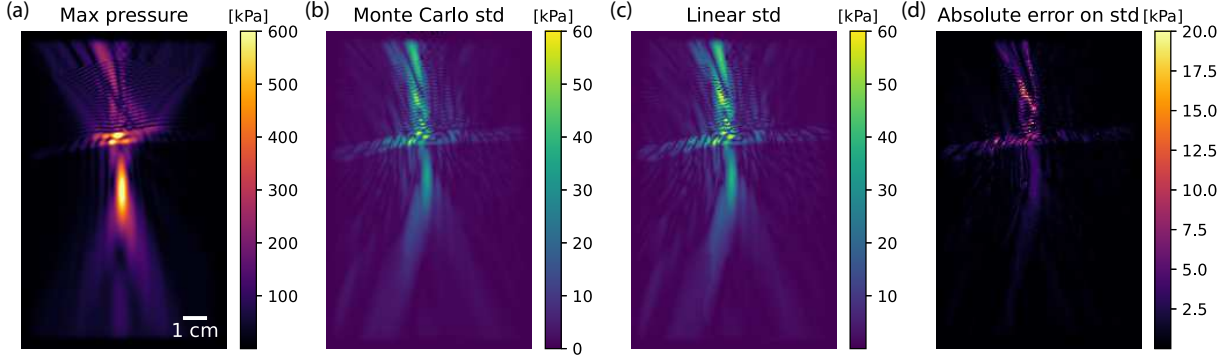


Figure 2: (a) The simulated pressure field, (b) the standard deviation (STD) of the simulation results using 200 Monte Carlo samples, (c) the standard deviation estimated by the linear uncertainty propagation method, and (d) the error between the estimated standard deviations.

Since the variables are uncorrelated, the formula for the variance simplifies to

$$\text{Var}(f(\mathbf{v})) = \left(\frac{\partial f}{\partial \alpha} \Big|_{\mathbf{v}=\mu_v} \right)^2 \sigma_\alpha^2 + \left(\frac{\partial f}{\partial \beta} \Big|_{\mathbf{v}=\mu_v} \right)^2 \sigma_\beta^2. \quad (6)$$

If the Jacobian is non-singular, this estimate can provide a good approximation for the variance for small variations of the parameter vector \mathbf{v} . To estimate the normal distribution of the outputs, all that is needed is an efficient way to calculate $f(\mu_v)$ and the Jacobian vector of f at $\mathbf{v} = \mu_v$.

3 Numerical Experiments

To demonstrate uncertainty estimation using Eq. (6), an ultrasound simulation through a CT-derived acoustic property map was performed in the context of TUS. A CT image of the head was taken from the dataset described in [14, 16], re-sampled to an isotropic resolution of $0.5 \times 0.5 \times 0.5$ mm, and truncated to a $120 \times 70 \times 70$ mm domain in the region of the primary motor cortex. The simulation uses speed of sound and density values derived from the CT image, while medium is considered lossless. The transducer was a focused bowl defined following [17], with a radius of curvature and aperture diameter of 64 mm, source pressure of 60 kPa, and driving frequency of 500 kHz.

Simulations were performed using j-Wave [13], a differentiable wave solver, by running a time-domain simulation to steady state, and then recording the maximum pressure amplitude over several cycles. To avoid staircasing artefacts, the transducer was represented on the grid using the band limited interpolant [18]. The grid spacing gave 6 points per wavelength (PPW) in water, and the time step was set using a Courant-Friedrichs-Lewy number of 0.3.

To calculate the gradients $\partial f / \partial \alpha$ and $\partial f / \partial \beta$, forward-mode differentiation was used [19, 13]. In j-Wave, this can be performed in parallel with the unperturbed forward computation, at a cost of one forward computation for each partial derivative. Forward-mode differentiation, in contrast to the widely-used adjoint method, does not necessitate storing the intermediate computations. In the context of wave physics it has been used, for instance, to perform sensitivity-analysis of photonic devices [20]. Alternatively, the gradients could also be approximated using a non-differentiable wave solver by replacing the gradients with a finite difference approximation.

As a reference, the standard deviation of the simulation was also estimated via Monte Carlo, by running 200 simulations with samples of the parameter vector \mathbf{v} (i.e., different values for the slope and intercept in the linear relationship between density and sound speed) drawn from its probability distribution.

4 Results

The simulated acoustic pressure field for a plane passing through the focus is shown in Fig. 2(a). The standard deviation calculated using Monte Carlo is displayed in Fig. 2(b), while the standard deviation calculated using linear uncertainty propagation is displayed in Fig. 2(c). The two uncertainty maps are in close agreement, both in terms of the structures shown and the size of the standard deviation estimate. The difference between the two, illustrated in Fig. 2(d), shows that the uncertainty is accurately estimated, particularly below the skull and close to the focal region, while the most noticeable discrepancies between the two are found inside the skull and in the reflected wave-field. The uncertainty

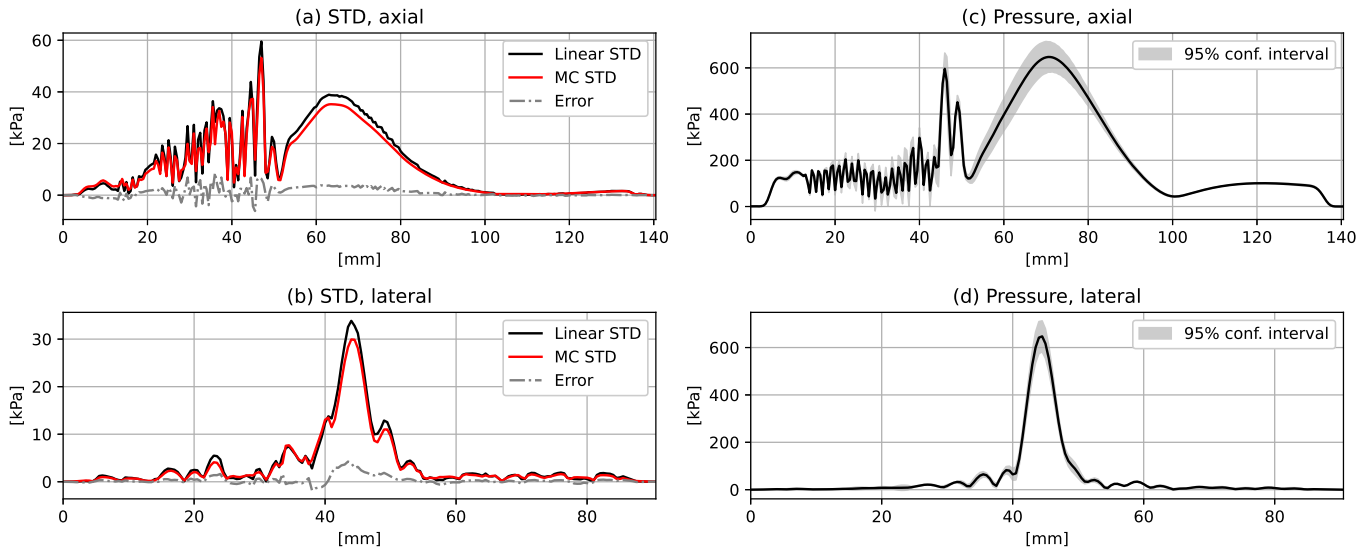


Figure 3: The standard deviations (STD) estimated using Monte Carlo (MC) and linear uncertainty propagation (Linear), and the difference error between the two, for the (a) axial and (b) lateral lines passing through the focus. Subplots (c) and (d) show the axial and lateral pressure profiles through the focus, along with the 95% confidence interval (2-sigma) estimated using linear uncertainty propagation.




Figure 4: The histogram of the empirical distribution of the pressure value at the focus, estimated using 200 Monte Carlo runs, and the corresponding normal distribution estimated by linear uncertainty propagation (LUP).

estimates for lateral and axial profiles through the focus are shown in Fig. 3, highlighting how the algorithm predicts an error of up to 40 kPa near the focal region, which is very close to the Monte Carlo estimate.

A representative example of the empirical probability distribution for the pressure value at the focus is shown in Fig. 4. This shows a histogram of the pressure values at a single grid point for each of the 200 Monte Carlo simulations. The normal distribution predicted by the linear uncertainty propagation is shown on top. Despite the fact that the empirical distribution appears to be slightly skewed and therefore is deviating marginally from the supposed Gaussian distribution, the approximated distribution still satisfactorily explains the spread of possible pressure values for the given pixel.

Figure 5 shows the convergence of the estimated uncertainty using Monte Carlo for 10 different random seeds. With small numbers of iterations (simulations), the Monte Carlo estimate has a large amount of variability. While the Monte Carlo solution will eventually converge to an unbiased value, this can take, on average, a very large number of simulations. In contrast, linear uncertainty propagation only requires two simulations, so is computationally efficient.




Figure 5: Example of 10 different Monte Carlo (MC) runs and their convergence in estimating the standard deviation at the focus, along with the uncertainty estimated using linear uncertainty propagation (LUP).

However, while for the simulation parameters used in this study it provides a good approximation of the true variance, for large values of ε , the biased nature of linear uncertainty propagation may start to become relevant.

5 Summary

We show that linear uncertainty propagation can be used to estimate uncertainty in simulated transcranial ultrasound fields when there is uncertainty in the medium property mapping between mass density and sound speed. This provides a computationally efficient way of putting error bars on acoustic simulations used for treatment planning due to uncertainty in the material properties. In future, we will explore accounting for additional uncertainties (for example, in the skull attenuation), the sensitivity of the predicted field to variations in different material properties, and evaluate the regime of validity of linear uncertainty propagation for different model parameters (e.g., transducer shape, frequency, target region, etc).

Another potentially interesting extension would be to directly couple the acoustic simulation with a thermal simulation to estimate the uncertainty on the temperature rise induced by the acoustic field (this is relevant for thermal ablation, as well as safety for non-thermal therapies). As long as the thermal simulation is also written using a differentiable language, the algorithm outlined in this paper will still be applicable.

Lastly, note that the algorithm derived for estimating uncertainty is itself differentiable. This allows the predicted uncertainty map to be included in an other optimization loop and the definition of a cost function that depends on uncertainty that is still optimizable using gradient descent, as done for example in [11] using a different method for estimating uncertainties.

Acknowledgments

The authors would like to thank Taylor Webb and Kim Butts Pauly for provision of the raw data from [6, 7] used to generate Fig. 1(b). This work was supported by the Engineering and Physical Sciences Research Council (EPSRC), UK, grant numbers EP/S026371/1 and EP/T022280/1.

References

- [1] T. Deffieux, E. E. Konofagou, Numerical study of a simple transcranial focused ultrasound system applied to blood-brain barrier opening, *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* 57 (12) (2010) 2637–2653.
- [2] N. McDannold, P. J. White, R. Cosgrove, Elementwise approach for simulating transcranial mri-guided focused ultrasound thermal ablation, *Physical review research* 1 (3) (2019) 033205.
- [3] W. Lee, H.-C. Kim, Y. Jung, Y. A. Chung, I.-U. Song, J.-H. Lee, S.-S. Yoo, Transcranial focused ultrasound stimulation of human primary visual cortex, *Scientific reports* 6 (1) (2016) 1–12.
- [4] N. Li, P. Gaur, K. Quah, K. Butts Pauly, Improving in situ acoustic intensity estimates using MR acoustic radiation force imaging in combination with multifrequency MR elastography, *Magnetic Resonance in Medicine* 88 (4) (2022) 1673–1689.

- [5] F. Marquet, M. Pernot, J.-F. Aubry, G. Montaldo, L. Marsac, M. Tanter, M. Fink, Non-invasive transcranial ultrasound therapy based on a 3D CT scan: Protocol validation and in vitro results, *Physics in Medicine & Biology* 54 (9) (2009) 2597.
- [6] T. D. Webb, S. A. Leung, J. Rosenberg, P. Ghanouni, J. J. Dahl, N. J. Pelc, K. B. Pauly, Measurements of the relationship between CT hounsfield units and acoustic velocity and how it changes with photon energy and reconstruction method, *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* 65 (7) (2018) 1111–1124.
- [7] T. D. Webb, S. A. Leung, P. Ghanouni, J. J. Dahl, N. J. Pelc, K. B. Pauly, Acoustic attenuation: Multifrequency measurement and relationship to CT and MR imaging, *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* 68 (5) (2020) 1532–1545.
- [8] T. E. Vaughan, K. Hynynen, Effects of parameter errors in the simulation of transcranial focused ultrasound, *Physics in Medicine & Biology* 47 (1) (2001) 37.
- [9] J. Robertson, E. Martin, B. Cox, B. E. Treeby, Sensitivity of simulated transcranial ultrasound fields to acoustic medium property maps, *Physics in Medicine & Biology* 62 (7) (2017) 2559.
- [10] H. Montanaro, C. Pasquinelli, H. J. Lee, H. Kim, H. R. Siebner, N. Kuster, A. Thielscher, E. Neufeld, The impact of CT image parameters and skull heterogeneity modeling on the accuracy of transcranial focused ultrasound simulations, *Journal of Neural Engineering* 18 (4) (2021) 046041.
- [11] A. R. Gerlach, A. Leonard, J. Rogers, C. Rackauckas, The Koopman expectation: An operator theoretic method for efficient analysis and optimization of uncertain hybrid dynamical systems, *arXiv preprint arXiv:2008.08737* (2020).
- [12] G. Papamakarios, I. Murray, Fast ε -free inference of simulation models with Bayesian conditional density estimation, *arXiv preprint arXiv:1605.06376* (2016).
- [13] A. Stanziola, S. R. Arridge, B. T. Cox, B. E. Treeby, j-Wave: An open-source differentiable wave simulator, *arXiv preprint arXiv:2207.01499* (2022).
- [14] J. Caballero-Insaurriaga, R. Rodríguez-Rojas, R. Martínez-Fernández, M. Del-Alamo, L. Díaz-Jiménez, M. Ávila, M. Martínez-Rodrigo, P. García-Polo, J. A. Pineda-Pardo, Zero TE MRI applications to transcranial MR-guided focused ultrasound: Patient screening and treatment efficiency estimation, *Journal of Magnetic Resonance Imaging* 50 (5) (2019) 1583–1592.
- [15] M. Giordano, Uncertainty propagation with functionally correlated quantities, *arXiv preprint arXiv:1610.08716* (2016).
- [16] M. Miscouridou, J. A. Pineda-Pardo, C. J. Stagg, B. E. Treeby, A. Stanziola, Classical and learned MR to pseudo-CT mappings for accurate transcranial ultrasound simulation, *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 69 (10) (2022) 2896–2905.
- [17] J.-F. Aubry, O. Bates, C. Boehm, K. Butts Pauly, D. Christensen, C. Cueto, P. Gelat, L. Guasch, J. Jaros, Y. Jing, R. Jones, N. Li, P. Marty, H. Montanaro, E. Neufeld, S. Pichardo, G. Pinton, A. Pulkkinen, A. Stanziola, A. Thielscher, B. Treeby, E. van ’t Wout, Benchmark problems for transcranial ultrasound simulation: Inter-comparison of compressional wave models, *The Journal of the Acoustical Society of America* 152 (2) (2022) 1003–1019.
- [18] E. S. Wise, B. Cox, J. Jaros, B. E. Treeby, Representing arbitrary acoustic source and sensor distributions in fourier collocation methods, *The Journal of the Acoustical Society of America* 146 (1) (2019) 278–288.
- [19] A. Stanziola, S. Arridge, B. T. Cox, B. E. Treeby, A research framework for writing differentiable PDE discretizations in JAX, *Differentiable Programming workshop at Neural Information Processing Systems 2021* (2021).
- [20] T. W. Hughes, I. A. Williamson, M. Minkov, S. Fan, Forward-mode differentiation of Maxwell’s equations, *ACS Photonics* 6 (11) (2019) 3010–3016.