JMI Journal of Medical Imaging

RESEARCH PAPER

Ultrasound elastic modulus reconstruction using a deep learning model trained with simulated data

Utsav Ratna Tuladhar,[a,*] Richard A. Simon,[b] Cristian A. Linte,[b] and Michael S. Richards[b]

[a]Rochester Institute of Technology, Electrical and Computer Engineering, Rochester, New York, United States
[b]Rochester Institute of Technology, Biomedical Engineering, Rochester, New York, United States

ABSTRACT. **Purpose:** Ultrasound (US) elastography is a technique for non-invasive quantification of material properties, such as stiffness, from ultrasound images of deforming tissue. The material properties are calculated by solving the inverse problem on the measured displacement field from the ultrasound images. The limitations of traditional inverse problem techniques in US elastography are either slow and computationally intensive (iterative techniques) or sensitive to measurement noise and dependent on full displacement field data (direct techniques). Thus, we develop and validate a deep learning approach for solving the inverse problem in US elastography. This involves recovering the spatial modulus distribution of the elastic modulus from one component of the US-measured displacement field.

Approach: We present a U-Net-based deep learning neural network to address the inverse problem in ultrasound elastography. This approach diverges from traditional methods by focusing on a data-driven model. The neural network is trained using data generated from a forward finite element model. This simulation incorporates variations in the displacement fields that correspond to the elastic modulus distribution, allowing the network to learn without the need for extensive real-world measurement data. The inverse problem of predicting the modulus spatial distribution from ultrasound-measured displacement fields is addressed using a trained neural network. The neural network is evaluated with mean squared error (MSE) and mean absolute percentage error (MAPE) metrics. To extend our model to practical purposes, we conduct phantom experiments and also apply our model to clinical data.

Results: Our simulated results indicate that our deep learning (DL) model effectively reconstructs modulus distributions, as evidenced by low MSE and MAPE evaluation metrics. We obtain a mean MAPE of 0.32% for a hard inclusion and 0.39% for a soft inclusion. Similarly, in our phantom studies, the predicted modulus ratio aligns with the expected range, affirming the model's accuracy. These findings, alongside evaluations using the modulus ratio and contrast-to-noise ratio, confirm our DL model's robust generalization capabilities across diverse datasets.

Conclusions: The presented work demonstrated that provided the simulated data are sufficiently diverse and representative of a wide variability, the algorithm trained on simulated data would generalize well to both phantom, as well as real-world clinical data.

 [DOI: 10.1117/1.JMI.12.1.017001]

Keywords: ultrasound elastography; inverse problem; deep learning

Paper 24158GRR received May 30, 2024; revised Dec. 23, 2024; accepted Jan. 10, 2025; published Feb. 5, 2025.

*Address all correspondence to Utsav Ratna Tuladhar, ut3320@rit.edu

2329-4302/2025/$28.00 © 2025 SPIE

1 Introduction

Ultrasound (US) imaging is a widely used clinical modality, known for its ease of use, real-time capability, portability, low cost, and low risk for patients. It uses acoustic impedance differences to image the morphological characteristics of soft tissues.[1] Ultrasound elastography (USE), or elasticity imaging, has become a powerful adjunct technique used for non-invasive assessment of tissue pathology by evaluating the mechanical properties of tissues.[2] USE applications include the detection of breast lesions,[3] assessment of liver fibrosis,[4,5] cardiovascular imaging,[6] and evaluation of thyroid nodules,[7] to name a few. Currently, shear wave–based USE is more commonly utilized in clinical settings due to its ability to directly relate shear wave speeds to the underlying shear modulus. However, due to the additional functionality required for shear wave generation, improving the accuracy of strain-based compression USE, particularly in real-time modulus imaging, could enhance its diagnostic applications in more portable, low-cost ultrasound systems. In compression USE, the body surface of the tissue of interest is subjected to a compressing force, often via the US transducer itself, which induces deformation within the tissue, an example of which is shown in Fig. 1. From ultrasound image sequences of this deformation, the displacement fields of the imaged tissues are measured. From the displacement fields, an axial strain field can be calculated and used to infer material properties.[8] However, it has been shown that strain alone is an inaccurate representation of material properties which motivates the need to recover the actual property distributions using measured displacement fields.[9] This leads to the elastic inverse problem, where estimates of the mechanical properties of tissue are the desired calculation from the input of the measured displacement field.[10]

Various strategies have been devised to tackle the complex inverse problem of estimating tissue mechanical properties from imaging data. These approaches can be broadly categorized into direct and iterative methods. A detailed review of these techniques can be found in Doyley.[11] Briefly, direct methods solve the inverse problem by rearranging the equations of motion and the constitutive equations, typically linear elastic models, to solve for the elastic modulus directly from the known displacement fields.[12,13] Although direct techniques can be computationally efficient, they are more susceptible to measurement noise[11] and may require full-field data (rather than sparse).[13] In contrast, iterative techniques use numerical solutions of the forward elasticity problem and guesses of the modulus values to predict displacement fields. The modulus guesses are then iteratively updated, typically guided by gradient or Hessian calculations of an optimization function until optimal convergence is achieved so that the model-predicted displacement field closely matches the observed data.[14] This iterative nature makes them adaptable to a wide range of scenarios, even when analytical solutions are unavailable. However, this flexibility comes at the cost of being computationally expensive, limiting their utility in real-time clinical applications. The balance of precision and computational efficiency in these methods remains a key challenge in the advancement of elastography techniques for practical clinical applications.[15]

In the past few years, significant progress development of deep learning algorithms has opened up new possibilities for establishing a mapping function between input and output images. With this advancement, deep learning has also begun to gain prominence in the field

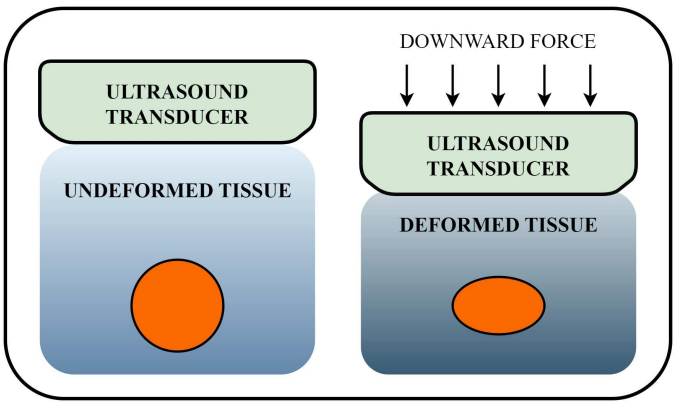

Fig. 1 Schematics of compression elastography. Here, the tissue gets perturbed by an external force applied by the US transducer.

of elasticity imaging.[16] In a quasi-static elastography paradigm, convolutional neural networks (CNNs) were first used to replace or improve tissue motion estimation techniques.[17–20] CNNs have also been used successfully to directly measure strain fields from radio frequency (RF) US images.[21] Gao et al.[22] trained an RF to strain CNN and incorporated tissue displacement as an intermediary step within the network. Traditional strain elastography and US B-mode images of breast and brain masses have also been used to train CNNs to directly classify and diagnose cancers.[23,24] Deep learning classification techniques have also been used with dynamic elastography for applications in breast tumors, liver imaging, and thyroid nodules.[16] Super-resolution techniques using generative adversarial networks (GANs) have also been applied to ultrasound elastography.[25] Although these works imply the effectiveness of the deep learning approach for ultrasound elastography, relying on displacement or strain images, rather than material property images, is still insufficient. In addition, avoiding image reconstruction, in favor of direct classification, must be a unique model developed for each specific pathology and cannot be used as a generalized imaging technique. Although direct classification may delineate healthy from unhealthy tissue, obfuscating the model's decision-making process can lead to further challenges.

Some preliminary studies have begun to develop deep learning techniques for the reconstruction of tissue material property distributions in USE; however, they have not yet been tested or implemented experimentally. Deep learning methods incorporating elasticity theory have been used to recover modulus distributions from observed displacement fields.[26,27] Here, a cost function is created that contains a neural network predicted elasticity distribution and measured strain fields to minimize the magnitude of the residual forces in the equilibrium equations. Then, training the neural net to minimize the cost function is synonymous with learning the elastic distribution. This technique is similar to the classical optimization framework for solving the inverse problem, without the need for calculating the gradients or Hessian's of the cost function. However, this technique may still suffer from similar computational limitations, as each measurement requires a new neural network to be trained. Mohammadi et al. also used a classical optimization framework with regularization, where the regularizer was trained using Wasserstein GAN to distinguish noisy from clean modulus images.[28] This technique worked effectively in simulated data but was not tested in real data. This technique also suffers from the high computational costs of iterative techniques.[28] Several groups have trained CNNs to quickly and efficiently map displacement or strain images to modulus images.[29–31] These neural networks were trained using simulated datasets, avoiding the need for clinical training databases, and once trained, they could be deployed on other data. Gokhale[29] compared CNNs trained with simulated datasets created of various complexities (1, 2, or 3 inclusions) and types (axial displacements versus axial strains). He found that using the axial displacement fields and adding complexity to the training set improved the accuracy of the predicted elasticity distributions. Ni and Gao[30] and Zhang et al.[31] independently introduced a discriminator network, rather than a mean squared error (MSE)-type cost function, to create conditional GANs (cGANs) to train their CNN. Although the cGAN architecture may prove more accurate than using a cost function, their models were trained on rather simplistic simulations with little complexity in the modulus images or boundary conditions.[30,31] In our implementation, we chose to start with an MSE cost function–based training and focus on the complexity of the trained dataset and its effect on the precision of the predicted modulus image.

Our previous work showed promising preliminary findings from phantom experiments, demonstrating the feasibility of using a model trained on simulated data to accurately predict the results of the phantom data.[32] In the work presented here, we substantially enhance the capabilities of our U-Net-based regression model. We first develop a novel and expansive dataset inclusive of diverse scenarios of displacement fields achieved through finite element (FE) forward modeling of corresponding modulus distribution. We then proceed to train a deep learning (DL) model with this dataset, using an MSE loss function. We demonstrate improved performance with this framework with respect to its ability to capture wider variability. In particular, we examine the influence of various region of interest (ROI) windows, which correspond to the various depths of the observations. We also conduct a noise study for the simulated data. Furthermore, to illustrate the practical applicability of our model in a clinical setting, we apply this DL model to phantom experiments featuring eight distinct inclusions that broaden our

experimental scope and show that our model can generalize from simulations to realistic conditions. We validate our phantom results using contrast and contrast-to-noise ratio (CNR) metrics. We also analyze the impact of variations in the ROI, demonstrating how our DL model performs with displacements obtained from ROIs captured at different locations within US images. Finally, we also incorporate the results from two clinical exams in our analysis, which allows us to evaluate the potential application of our DL model in a real-world clinical context. In summary, the contributions of this study are as follows: (1) development of a comprehensive dataset, (2) DL model enhancement and ROI analysis, and (3) application and analysis of our model on phantom experiments and clinical case studies.

2 Data

2.1 Simulated Data with Finite Element Modeling

To learn the relationship between displacement fields and their corresponding modulus distributions using a data-driven deep learning approach, we created a comprehensive dataset that covers all possible variations. To that end, we simulated data by solving two-dimensional (2D) plane stress forward problems, that is, calculating the displacement field from the modulus distribution, using the FE method (FEM) with an internal algorithm. The inputs for the FE model were the spatial distribution of the 2D modulus, the geometry of the tissue sample discretized as an FE mesh, and the boundary conditions. The FE mesh consists of evenly distributed nodes on a homogeneous 128×128 Cartesian grid connected by quadrilateral elements. The modulus and displacement fields are discretized at the nodal locations. With these inputs, the FE model outputs the displacement field in the x and y directions. Detailed simulation methodology has been described by Tuladhar et al.[32]

We generate the modulus distribution by adjusting the shape and size of each inclusion, starting with a circle and varying its radius based on a sum of cosine functions with set amplitudes and phases,[33] using the equation

$$r(\theta) = A\left(1 + \sum_{n=1}^{5} B_n \cos(2\pi n\theta + \psi)\right), \tag{1}$$

where B_n is the amplitude that varies between $0.25 \leq B_n \leq 0.5$, ψ is the phase with $0 \leq \psi \leq 2\pi$, and A is the scaling factor with $0.113L \leq A \leq 0.225L$. We also randomly varied the locations of these inclusions within the distribution space. A Gaussian filter with a standard deviation uniformly ranging from 0.5 to 4 was applied to the modulus distribution to smooth out sharp edges. Variations in boundary conditions are described below. At $y = 0$, $u_x = u_y = 0.0$ as the bottom is fixed. Similarly, at $y = L$, we have $-0.005 \leq u_x \leq 0.005$ to account for the lateral shift of the ultrasound probe due to human error in otherwise perfect vertical compression. Also, at $y = L$, we have $u_y = -0.01L + mx + b$, where b and x_0 are chosen from a uniform random distribution within the range $-0.005 \leq b \leq 0.005$ and $0.33L \leq x_0 \leq 0.66L$, respectively, from which we also calculate $m = b/x0$. In addition, the edges were set to be traction-free. A plane stress formulation, assuming a Poisson ratio of 0.49, was used to model the material as a linear elastic isotropic material. In this work, we use only the axial displacement field (Y direction), u_y, to train our deep learning model.

In addition to the variations described in our previous work, we add another set of variations to the domain of the Cartesian grid by increasing the domain of the FE mesh from 128×128 to 200×200 and 400×400 which mimics the possibility of inclusion being in a broader area of the human body. The variability in the domain mimics the variability in the size of the body tissue within which the inclusion is located. We get outputs in the 200×200 and 400×400 grids, from which we extract a randomly selected window of 128×128 to train our model. The greater domain size and selection of smaller windows from these domains also provide us with more set of variability in the boundary condition. In this way, we created a dataset that captures wider variations of displacement fields. Moreover, these new variations in boundary conditions also incorporate edge cases, where the inclusion is not well centered in the imaged data.

In this forward modeling, as shown in Fig. 2, we keep several parameters constant. First, the displacement boundary condition at $y = 0$ signifies that along the bottom boundary of the

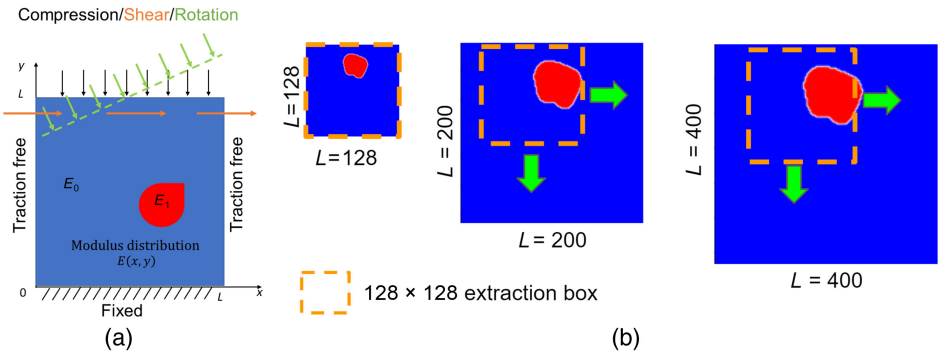

Fig. 2 FEM simulation framework. (a) FEM input model illustrating the input parameters for the finite element model, highlighting both the variable and fixed parameters. (b) Domain variation schema with variable domain size (L) demonstrating the variation of the ROI across different domains, showcasing the method for adjusting ROI in the simulation.

structure, there are no horizontal and vertical movements allowed. The Poisson ratio was set to 0.49, which characterizes the material's incompressible behavior under deformation. The background modulus (E_0) is set at a fixed value of 1.0, which represents the nominal stiffness of the background of the material. These constants serve as essential assumptions in the analysis and help simplify the modeling of the structure under specific conditions. 2D plane stress was also assumed.

The variable elements are shown graphically in Fig. 2. We used boundary conditions to mimic tissue compression from the ultrasound probe. The displacement boundary condition at $y = L$ has variations in compression (black), shear (yellow), and rotation (green). These represent the variations that might be realized by different technicians when applying compression. The shape and size of inclusion were also varied with the help of enclosed circular areas (red). A Gaussian filter, with a standard distribution uniformly between 0.5 and 4, was then applied to this modulus distribution to eliminate sharp edges. We also randomly vary the location of these inclusions within the distribution space. We note here that, as it is not possible to measure the force applied by ultrasound boundary, the elastic modulus can only be reconstructed relative to the elastic modulus of the background.

Our training set consists of 30,000 examples, 10,000 each from the 128, 200, and 400 domains. A total of 250,000 examples of 128×128-sized image pairs were extracted featuring both hard inclusions and soft inclusions. Similarly, we created a validation and test dataset containing 25,000 examples each.

We also investigate the impact of noise on the simulated data by adding Gaussian noise to the displacement fields prior to normalization. Gaussian noise was added with different values of standard deviations to emulate different noise levels. The amount of noise was quantified by a strain amplitude, signal-to-noise ratio (SNR_e) using strains calculated from our simulated displacement fields.[34,35] The SNR_e is given by the equation

$$\mathrm{SNR}_e = \frac{\sqrt{\frac{1}{N_{\mathrm{pixels}}}\sum_{i=1}^{N_{\mathrm{pixels}}} \epsilon_i^2}}{\sqrt{\frac{2\sigma_n^2}{W^2}}}, \tag{2}$$

where ϵ is the strain, σ_n is the noise level, N_{pixels} is the number of pixels, and W is the pixel spacing. The denominator of Eq. (2) is an approximation to the standard deviation of the noise in the strain estimates, given the standard deviation of the displacement noise, when a finite difference calculation is used to estimate the strain. We used a value of $W = 0.3$ mm, corresponding to our pixel spacing, and σ_n values between 1×10^{-4} and 1×10^{-3} mm.

2.2 Phantom Data

To ensure the robustness and precision of our technique before applying it to patient data, we performed phantom experiments. These experiments provide a controlled environment that closely mimics real tissues, allowing us to test our method in a predictable medical setting.

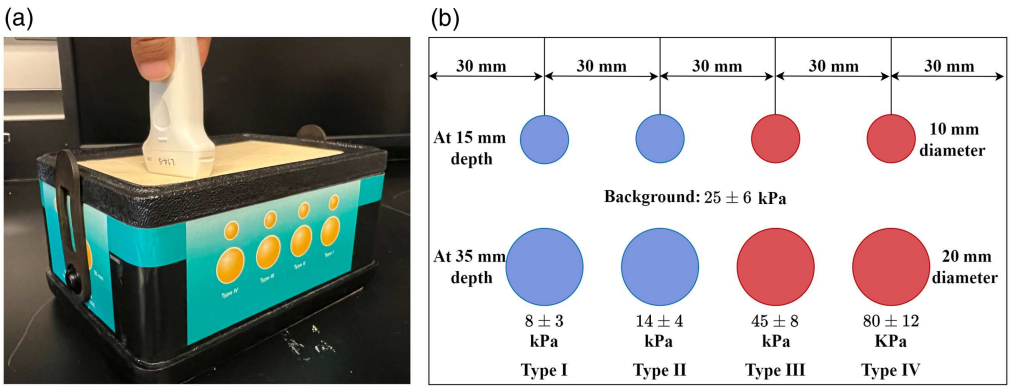

Fig. 3 Phantom structure overview. (a) Image of the physical phantom setup. (b) Detailed schematic representation of the phantom, indicating the precise locations, dimensions, and values of each inclusion within it.

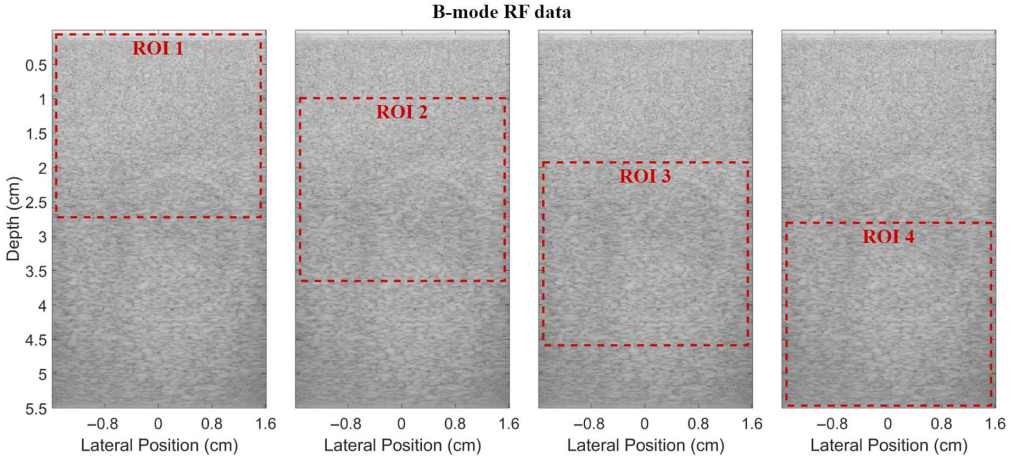

Fig. 4 Phantom ROI selection: B-mode images showing four extracted ROIs from various regions.

Here, we used a tissue-mimicking elasticity phantom, Elasticity QA Ultrasound Phantom CIRS 049, to collect data that contain both hard and soft spheres.[32] This phantom has a total of eight inclusions, each with specific values of elastic modulus as shown in Fig. 3. The elastic modulus of the background material is 25 ± 6 kPa. Dividing each modulus value by the background gives us the expected contrast and the expected contrast range, as shown in Table 3. These inclusions help simulate various tissue stiffness and size.

The experimental procedure we carried out for the phantom data is as follows. We collected phantom RF data using a Sonix Touch ultrasound scanner equipped with an L14-5/38 linear probe (Ultrasonix Medical Corporation, Richmond, British Columbia, Canada). We used standard beamforming with the focal point strategically placed in the inclusion area. With the US probe, we manually applied a slight axial compression (5% to 10%) during the collection process, and the images were acquired at a frame rate of ~123 Hz. Each RF image frame had dimensions of 2864×128 samples, equivalent to a physical area of 55 mm by 38 mm, with 128 RF signal scan lines sampled at 40 MHz. The imaging pulse was centered at 10 MHz. The phantom housed eight distinct inclusions, each subjected to five repeated measurements. For each measurement, four specific ROIs were analyzed. The RF data were further processed with the image registration algorithm to obtain the displacement field. From these displacement fields, ROIs were selected from various regions as shown in Fig. 4.

2.3 Clinical Data

To illustrate the practicality of our method in clinical settings, we employ two patient datasets from a study conducted by Rivaz et al.[36] This dataset includes RF data collected from patients

undergoing open surgical RF thermal ablation for liver cancer treatment between February 6, 2008, and July 28, 2009. This study included patients from John Hopkins Medicine, Baltimore, Maryland, United States. Patients enrolled in the study had diseases that could not be surgically removed and were selected for RF ablation after review. Patients with cirrhosis or suboptimal tumor location were not included in the study. US images were acquired after ablation for the two patients. Using these RF data, we first measure the displacement fields using the image registration algorithm. Consequently, we apply our neural network model to obtain modulus distribution results. This application of our method to clinical data underscores its relevance and potential utility in real-world medical scenarios.

3 Methods

3.1 Image Registration for Displacement Estimation

The axial displacement fields for the phantom and clinical RF US images were measured using a custom image registration algorithm written in MATLAB (MathWorks, Natick, Massachusetts, United States). The input to the algorithm is two frames, one reference image, and one target image from the US sequence, and the output is the measured 2D displacement field within the measurement domain (Ω). The basic framework of the registration algorithm can be found in our previous work.[34,37] The previous algorithm was adapted here to utilize a novel mechanics–based spatial regularization to strongly penalize measurements inconsistent with plane stress deformations. The CNN model was exclusively trained with materials that deform under plane stress conditions. Thus, we chose to develop a mechanics-based regularizer to minimize inconsistencies between our simulated training data and experimental data.

The optimization function used here was

$$\pi[\boldsymbol{u}(\boldsymbol{x})] = \frac{1}{2}\int_{\Omega}(I_1(\boldsymbol{x}) - I_2(\boldsymbol{x}+\boldsymbol{u}(\boldsymbol{x})))^2 \mathrm{d}\Omega + \alpha \int_{\Omega} |\nabla \cdot \boldsymbol{A}(\boldsymbol{x})| \mathrm{d}\Omega. \tag{3}$$

The first term in this function is the squared difference among the undeformed, reference image $(I_1(\boldsymbol{x}))$ and a target image $(I_2(\boldsymbol{x}))$, which is artificially warped, non-linearly, by the 2D displacement vector field $(\boldsymbol{u}(\boldsymbol{x}))$. The second term is a momentum equation-based regularization and is similar to that used in Babaniyi et al.[38] to post-process cross-correlation–based displacement estimates. In Eq. (3), $\boldsymbol{A}(\boldsymbol{x})$ is the constitutive relation for a linear elastic, homogeneous $(\mu(\boldsymbol{x}) = 1)$, incompressible material deforming under plane stress and can be written as[38]

$$\boldsymbol{A}(\boldsymbol{x}) = 2(\nabla \cdot \boldsymbol{u}(\boldsymbol{x}))\mathbf{I} + \nabla \boldsymbol{u} + \nabla \boldsymbol{u}^T. \tag{4}$$

$\nabla \cdot \boldsymbol{A}(\boldsymbol{x})$, from Eq. (3), is the static equilibrium equation for such a material (Babaniyi et al. formulated $\boldsymbol{A}(\boldsymbol{x})$ with respect to the strains, $\boldsymbol{\epsilon} = \frac{1}{2}(\nabla \boldsymbol{u}^T + \nabla \boldsymbol{u})$, and $\boldsymbol{\epsilon}(\boldsymbol{x})$ and $\boldsymbol{u}(\boldsymbol{x})$ were both discretized using bilinear FE shape functions. An additional compatibility term, also imposed as a penalty, was used to couple the strain and displacement estimates.)[38] In this work, $\boldsymbol{u}(\boldsymbol{x})$ was discretized using bilinear, quadrilateral FE shape functions which yields a piece-wise constant approximation (within elements) to the constitutive relation defined in Eq. (4). We further assumed that for a piece-wise constant $\boldsymbol{A}(\boldsymbol{x})$, the regularization term can be approximated as

$$\alpha \int_{\Omega} |\nabla \cdot \boldsymbol{A}(\boldsymbol{x})| \mathrm{d}\Omega \approx \alpha \sum_{k=1}^{K_e} \beta_k |\Delta \boldsymbol{A}_k \cdot \boldsymbol{n}_k|^2, \text{ where} \tag{5}$$

$$\beta_k = \frac{L_k}{\sqrt{|\Delta \boldsymbol{A}_k \cdot \boldsymbol{n}_k|^2 + \delta}}. \tag{6}$$

In these equations, K_e is the number of element edges within the interior of the FE domain, $\boldsymbol{n}_k$ is the unit vector normal to the k'th edge and L_k is the edge length. The tensor $\Delta \boldsymbol{A}_k$ is the difference in $\boldsymbol{A}$ among the elements that share the k'th edge, evaluated at the midpoint of the element edge. The δ parameter is a small constant ($\delta = 1\mathrm{e}-8$) to avoid a singularity in the gradient of the absolute value function when $\Delta \boldsymbol{A}_k = 0$. To simplify the calculation of the gradient of Eq. (3), β_k is calculated with the current guess of the displacement field but held fixed at each iteration. The regularization parameter value was subjectively chosen, via visual inspection of

several example axial strain fields, to minimize the noise without over-smoothing the solution. A value of $\alpha = 1e-4$ was used for all the phantom, and $\alpha = 1e-3$ was used for the clinical measurements found in this work. Input frames were chosen for each measurement such as there was ~1% total axial strain between the reference and target images. The functional minimization was performed using a Gauss–Newton method described previously and output a 2D displacement vector field.[34,39] The axial component of the measured displacement field ($u_y(\boldsymbol{x})$) was then interpolated onto a regular grid within the appropriate ROI as input to the deep learning model to obtain the modulus predictions.

3.2 Deep Learning for Modulus Reconstruction

We employ a deep learning architecture inspired by U-Net[40] to model the relationship between displacement fields and the modulus distribution, as illustrated in Fig. 5. This approach draws inspiration from the Pix-to-Pix framework[41] commonly used for image-to-image translation. However, unlike the conventional Pix-to-Pix setup that utilizes both generator and discriminator components, our approach opts for the exclusive use of the generator component and a loss function instead.

In this architecture, we used seven downsampling layers in the encoder block and seven upsampling layers in the decoder block. The latent bottleneck block situated between the encoder and decoder has a dimensionality of 512, thereby capturing and learning 512 distinctive features from the input displacement field to derive modulus distribution. Furthermore, we employ the MSE loss function for optimization.[32] The MSE loss is computed as follows:

$$\text{MSE Loss} = \frac{1}{N_{\text{pixels}}} \sum_{i=1}^{N_{\text{pixels}}} (E_i^t - E_i^p)^2, \tag{7}$$

where N_{pixels} is the number of pixels to calculate the loss function between the predicted E_i^p and the ground truth E_i^t modulus.

For the evaluation of our deep learning model accuracy, we calculated the mean absolute percentage error (MAPE) metric

$$\text{MAPE} = \frac{1}{N_{\text{pixels}}} \sum_{i=1}^{N_{\text{pixels}}} \left| \frac{E_i^t - E_i^p}{E_i^t} \right|, \tag{8}$$

which returns a normalized metric of modulus accuracy. Furthermore, we generate mask images to separate inclusion from the background using Otsu's thresholding algorithm.[42] With the binarized mask images of the predicted B_i^p and the ground truth B_i^t modulus, we calculate the Dice score using

$$\text{Dice score} = \sum_{i=1}^{N_{\text{pixels}}} \frac{2B_i^t B_i^p}{B_i^t + B_i^p}. \tag{9}$$

Using this mask, we also compute the relative modulus error (RME), the normalized error between the mean ground truth and prediction inclusion values, given by the equation

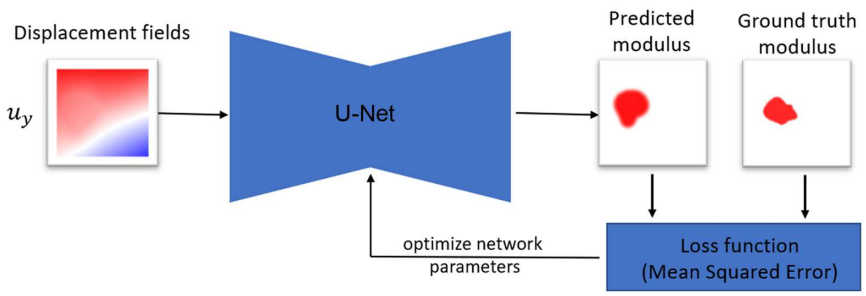

Fig. 5 Deep learning U-Net model featuring an MSE loss function for optimization.[32]

$$\mathrm{RME} = \frac{|\overline{E_{\mathrm{in}}^{t}} - \overline{E_{\mathrm{in}}^{p}}|}{\overline{E_{\mathrm{in}}^{t}}}, \tag{10}$$

where $\overline{E_{\mathrm{in}}^{t}}$ and $\overline{E_{\mathrm{in}}^{p}}$ are the mean values of ground truth modulus and predicted modulus respectively within the inclusion area.

To simulate the various situations that a technician might encounter when collecting data from ultrasound imaging, we organized the data by extracting samples from images of various domains. Our dataset includes a total of 250,000 training images and 25,000 test images. Specifically, we started with 10,000 samples directly taken from the 128×128 image domain. To augment this, we extracted 80,000 additional samples, each of size 128×128 pixels, from a collection of 10,000 images from the 200×200 image domain, effectively generating eight distinct samples from each image. Furthermore, from another set of 10,000 images of the 400×400 image domain, we derived 160,000 samples, also of size 128×128 pixels, by creating 16 unique samples from each original image. For data pre-processing, we normalized the displacement fields to values between 0 and 1. We trained our model using TensorFlow in our computing cluster.[43] We used a batch size of 128, with 200 epochs, and an Adam optimizer with a learning rate of 2×10^{-4}.

During the training process, we used the validation set to monitor the model's performance, assessing it at regular intervals. The model with the lowest MSE score on the validation set at the 159th epoch (Fig. S1 in the Supplementary Material) was saved and used for further evaluation. Following the completion of our DL model's training, we proceeded to apply this trained model for predicting new displacements obtained from the simulated test data, phantom data, and clinical data to evaluate our model's performance.

4 Results

To validate the performance of our model, we use data from our simulated test set, phantom experiments, and clinical trials. The results were obtained from the trained DL model as described in Sec. 3.2. Moreover, for phantom and clinical data, we used an image registration algorithm to measure the displacement fields of the ultrasound images.

4.1 Simulated Test Data

The test dataset was used to evaluate our DL model in simulations. To estimate the elastic modulus reconstruction error, we calculated the MSE and MAPE between ground truth and the predicted image. The average MSE obtained was 2.3×10^{-4} arbitrary units (a.u.), and the average MAPE was 0.33% for all test data. The processing time was on average 0.15 s per example in the test set. The average error observed for the hard inclusions, the soft inclusions, and the homogeneous modulus distributions is summarized in Table 1.

The validation of our DL model on the test dataset helped us assess its effectiveness in modulus reconstruction. The results, particularly the low MSE, signify that the predicted values are, on average, very close to the actual values in the dataset. Similarly, the low MAPE further confirms the model's accuracy in relative terms, indicating that the predictions are not only close in absolute terms but also in their deviation from the ground truth as a percentage. Furthermore, the results for both hard and soft inclusions within the dataset help us to understand the model's performance across different material types. Concurrently, the high Dice scores reflect a

Table 1 Summary of the mean and standard deviation values for ground truth and reconstructed elastic modulus for hard and soft inclusions. Comparison metrics include the mean MAPE, mean RME, and mean Dice score.

Inclusion	No. of images	Mean MAPE (%)	Mean RME (%)	Mean DICE
Hard	11,119	0.32 ± 0.12	1.38 ± 2.81	0.97 ± 0.12
Soft	10,751	0.39 ± 0.28	2.66 ± 9.15	0.99 ± 0.09
Homogeneous	3130	0.15 ± 0.01		—

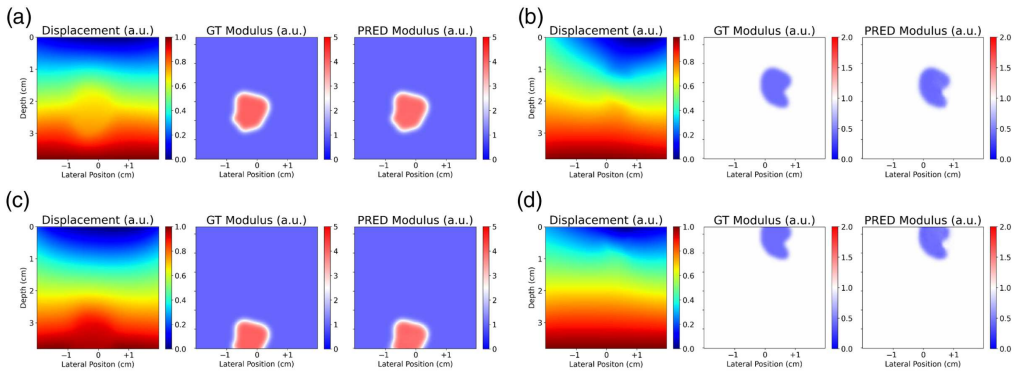

Fig. 6 Modulus distribution analysis in simulated data for different inclusions: Depicts the ground truth and predicted modulus for (a) hard inclusion situated within the ROI and (b) soft inclusion within the ROI. Modulus distribution for a (c) hard inclusion at the ROI's boundary and (d) soft inclusion at the ROI's boundary. The pixel size for all the images here is 128 by 128.

substantial overlap between the predicted and actual locations of the modulus, demonstrating the model's ability to identify the correct spatial positioning of the inclusions in the simulated data.

4.1.1 *ROI variations*

Examples of the predicted modulus can be seen in Fig. 6. An example of a hard inclusion is shown in Fig. 6(a). For this prediction, the MSE was 13.2×10^{-5}, and the MAPE was 0.307%. The average value of the ground truth modulus is 3.65, and the predicted value was found to be 3.67 with a relative error of 0.401% and a Dice score of 0.998. An example of a soft inclusion is shown in Fig. 6(b). For this prediction, the MSE was 1.58×10^{-5}, and the MAPE was 0.391%. The ground truth average modulus value was found to be 0.377, and the predicted value was 0.383 with a relative error of 1.664% and a dice score of 0.998. An example of a hard inclusion with an edge case is shown in Fig. 6(c). For this prediction, MSE was 6.69×10^{-5}, and MAPE was 0.249%. The ground truth average modulus value was found to be 2.025, and the predicted was 2.005 with a relative error of 0.981% and a Dice score of 0.994. An example of a soft inclusion with an edge case is shown in Fig. 6(d). For this prediction, the MSE was 7.13×10^{-5}, and the MAPE was 0.444%. The ground truth average modulus value was found to be 0.476, and the predicted value was 0.487 with a relative error of 2.381% and a Dice score of 0.999.

Our DL model has shown promising performance in handling a broader range of boundary conditions. These results demonstrate its effectiveness both when the entire object of interest is completely within the ROI and when only a portion of the object falls within the ROI. This versatility in handling different scenarios showcases the robustness and flexibility of our DL model in various real-world applications.

4.1.2 *Comparison to prior model*

To demonstrate the advantage over our previous model, Fig. 7 shows the predictions of our current DL modulus reconstruction with reconstruction results from Tuladhar et al.[32] The observed disparity in the performance between the current and previous DL models is notably attributed to the latter's limited capacity to handle variations, particularly evident in challenging edge cases. This deficiency is demonstrated in Fig. 7, where the current model exhibits more adaptability. The evidence presented here underscores robustness in model architecture and training, affirming the efficacy of our deep learning framework in handling complex edge cases.

4.1.3 *Noise study*

We evaluated the performance of our model by systematically increasing the noise and calculating the MSE and MAPE metrics of the output modulus images. The results of this study are

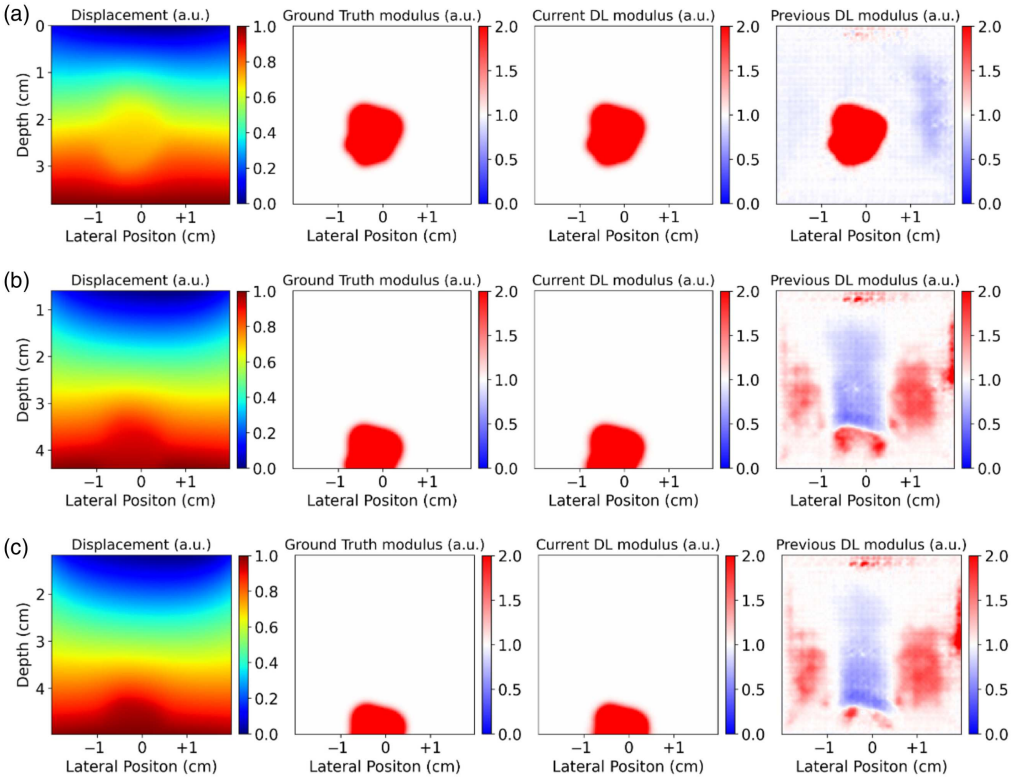

Fig. 7 Comparison between our previous and currently updated DL models for simulated test data. (a) Qualitatively better prediction when inclusion is completely inside. (b) Prediction getting worse as inclusion approach edge. (c) Prediction is completely missed at the edges.

Table 2 Summary of the mean and standard deviation SNR_e values for different noise levels on the test set, along with the mean and standard deviation MAPE values, illustrating the impact of noise on our model's performance.

Input: SNR_e	2.73 ± 1.05	0.91 ± 0.35	0.55 ± 0.21	0.34 ± 0.13	0.27 ± 0.10
Output: MAPE (%)	2.04 ± 1.20	11.67 ± 5.85	23.53 ± 10.49	42.81 ± 17.45	56.26 ± 22.16

summarized in Table 2. An example of a hard inclusion is shown in Fig. 8 using a sample from the simulated dataset. The MAPEs obtained are 0.43%, 1.48%, 3.93%, 10.92%, and 13.73% for SNR_e values of 5.83, 1.94, 1.17, 0.73, and 0.58, respectively.

Our analysis showed that as the noise level increased, the performance of the model to predict the modulus distribution degraded, as shown by the increased MAPE metric as the noise standard deviation increased.

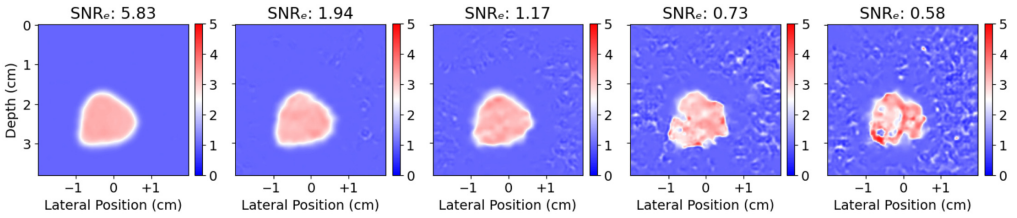

Fig. 8 Example modulus image predictions, from a hard inclusion, reconstructed from simulated displacement images with various noise levels added.

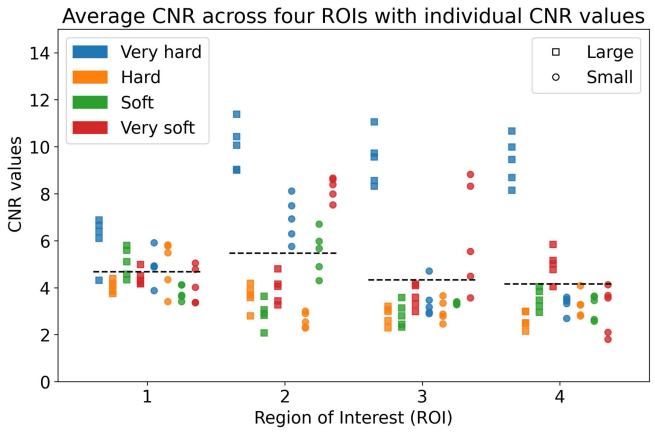

Fig. 9 Scatter plot depicting CNR values across different ROIs for various inclusion types. The dashed lines represent the mean values.

4.2 Phantom Data

After measuring the displacement fields from the phantom, we input them into our DL model to obtain the modulus distribution. To evaluate our phantom results, we focus on assessing the contrast in our DL output images as a measure of ground truth with the help of modular ratio. Modular ratio, in this context, refers to the degree of difference that makes an object distinguishable from the background.[44] To calculate this ratio we binarize the DL output using Otsu's thresholding algorithm[42] to separate the foreground and the background. Taking the ratio of inclusion's modulus to the background, we get the modular ratio by the equation

$$\text{Modular Ratio}(\eta) = \frac{\text{Modulus of inclusion}}{\text{Modulus of background}}. \tag{11}$$

We also calculate CNR along with contrast to analyze the delineation quality of inclusions in the reconstructions, given by

$$\text{CNR} = \sqrt{\frac{2(\hat{s}_t - \hat{s}_p)^2}{\hat{\sigma}_t^2 + \hat{\sigma}_p^2}}, \tag{12}$$

where $\hat{s}_t, \hat{s}_p$ are the mean, and $\hat{\sigma}_t, \hat{\sigma}_p$ are the standard deviation of the inclusion and background, respectively.

We compute the modular ratio and the CNR for each scenario. The CNR values are presented in a scatter plot in Fig. 9. Analysis of these CNR values indicates that ROI 2 yields better overall results. Consequently, the predicted contrast values for ROI 2 are tabulated in Table 3.

Table 3 Summary of the mean and standard deviation values for reconstructed modulus for various types of phantom inclusions for ROI 2.

Inclusion	Size	Expected η range	Expected η	Predicted η	CNR
Very hard (type IV)	Large	2.19 to 4.84	3.2	3.86 ± 0.04	10.01 ± 4.50
	Small	2.19 to 4.84	3.2	3.05 ± 0.15	4.30 ± 3.60
Hard (type III)	Large	1.19 to 2.79	1.8	1.61 ± 0.02	3.65 ± 31.89
	Small	1.19 to 2.79	1.8	1.45 ± 0.12	2.62 ± 1.32
Soft (type II)	Large	0.32 to 0.95	0.56	0.65 ± 0.06	2.94 ± 1.78
	Small	0.32 to 0.95	0.56	0.54 ± 0.01	5.57 ± 3.21
Very soft (type I)	Large	0.16 to 0.58	0.32	0.46 ± 0.03	3.94 ± 2.22
	Small	0.16 to 0.58	0.32	0.41 ± 0.01	8.24 ± 2.78

From this table, we can see that the average predicted contrast is well within the range of expected contrast. By conducting these phantom experiments, we establish a foundation upon which we can confidently proceed to apply and test our method on actual patient data.

4.2.1 *ROI variations*

Figure 10 illustrates the measured axial displacement field, the corresponding strain, and the predicted modulus for the type IV hard sphere. It presents the DL predictions for ROIs 1, 2, and 3, demonstrating the depth-wise variation. The predicted contrast values are 3.91, 3.87, and 3.88, respectively, and the CNR values are 6.66, 11.38, and 11.05. Similarly, Fig. 11 shows predictions for type I soft spheres. It presents the DL predictions for ROIs 1, 2, and 3, with the predicted contrast values 0.25, 0.40, and 0.39, respectively, and the CNR values are 4.78, 8.60, and 8.32.

The figures indicate that with ROIs that include the ultrasound transducer point of contact within the displacement fields, more noise or artifacts are affecting the predictions. With other ROIs such noise is comparatively reduced, leading to clearer predictions. The time taken for the DL model to reconstruct one displacement field was 0.15 s on average.

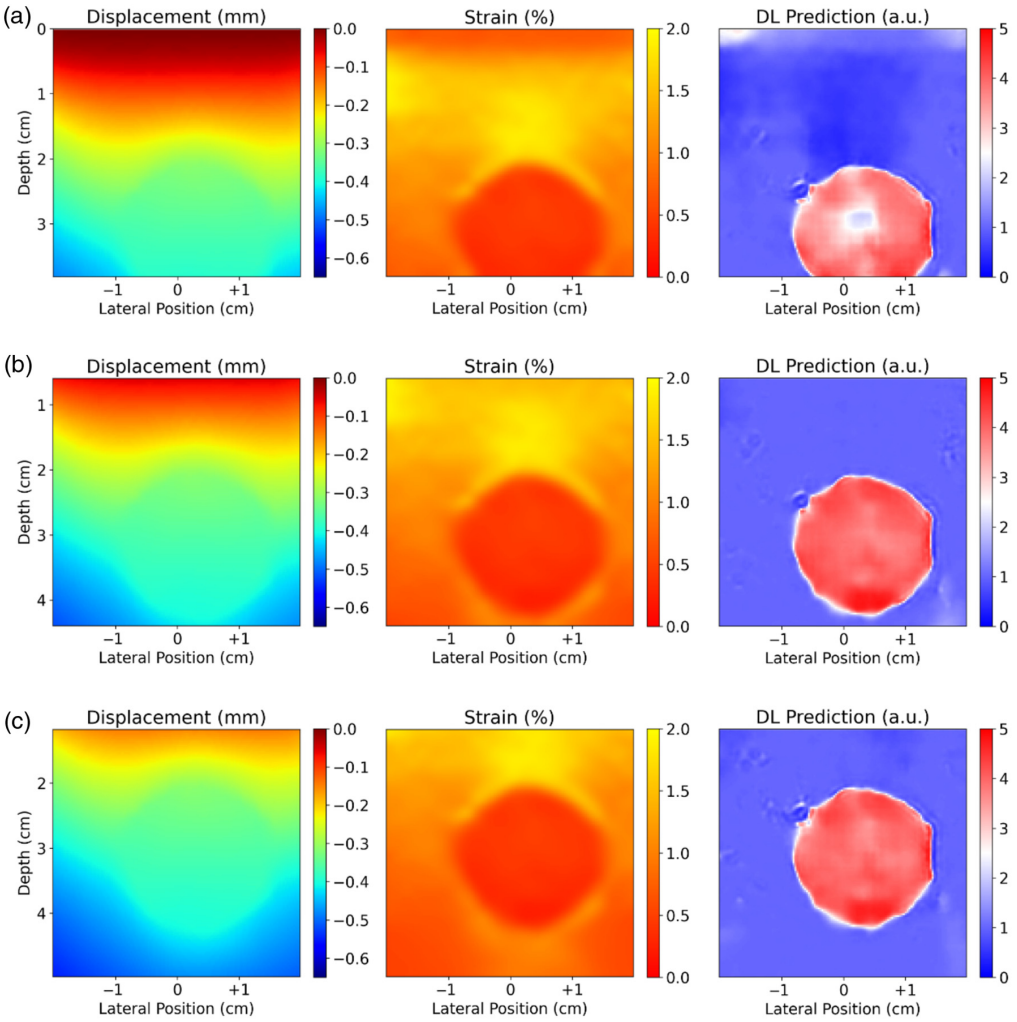

Fig. 10 Very hard inclusion from phantom with variability and ROIs showing reconstruction when inclusion lies (a) at the edge of the ROI, (b) near the edge of the ROI, and (c) completely inside the ROI. Every subfigure includes input displacement fields fed into the DL model, followed by the strain (1% average strain) and the modulus reconstruction.

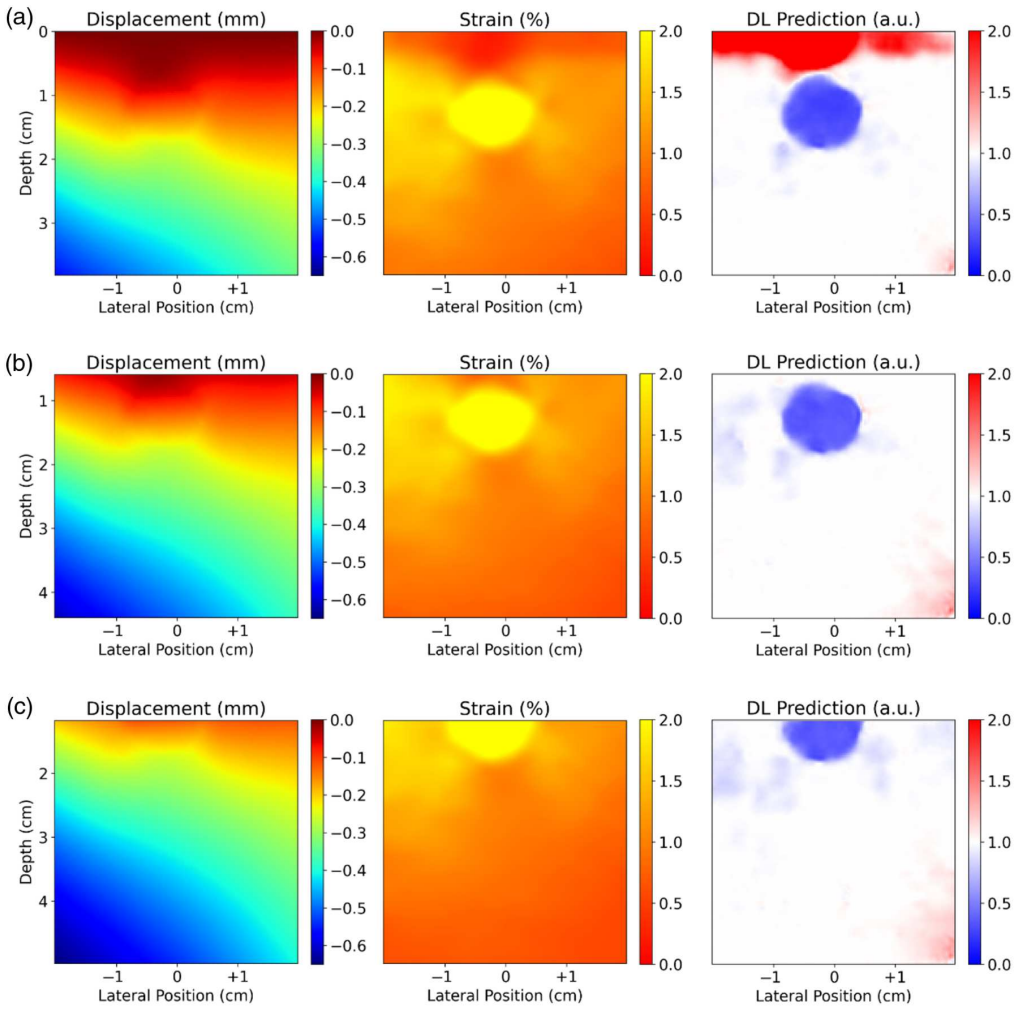

Fig. 11 Very soft inclusion from phantom with variability and ROIs showing reconstruction when inclusion lies (a) completely inside the ROI, (b) near the edge of the ROI, and (c) at the edge of the ROI. Every subfigure includes input displacement fields fed into the DL model, followed by the strain (1% average strain) and the modulus reconstruction.

In this study, we conducted five trials to evaluate the predictive performance and reliability of our deep learning model, as reflected in Fig. 12. The figure illustrates the model's response to a type II soft inclusion, with an anticipated contrast of 0.56, and an observed contrast range from 0.32 to 0.95. The obtained contrasts for trials 1 through 5 were 0.63, 0.62, 0.57, 0.61, and 0.65, demonstrating the model's consistent predictive capability. Correspondingly, the CNR recorded were 1.52, 1.56, 1.89, 1.67, and 1.77, further evidencing the model's robustness and reliability in varied measurement scenarios. This approach underscores our model's ability to maintain prediction accuracy and stability across multiple evaluations.

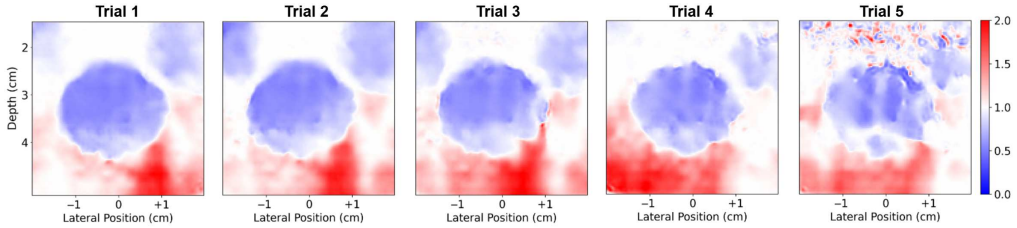

Fig. 12 Predictions of repeated measurements for type II soft inclusion phantom data from five trials.

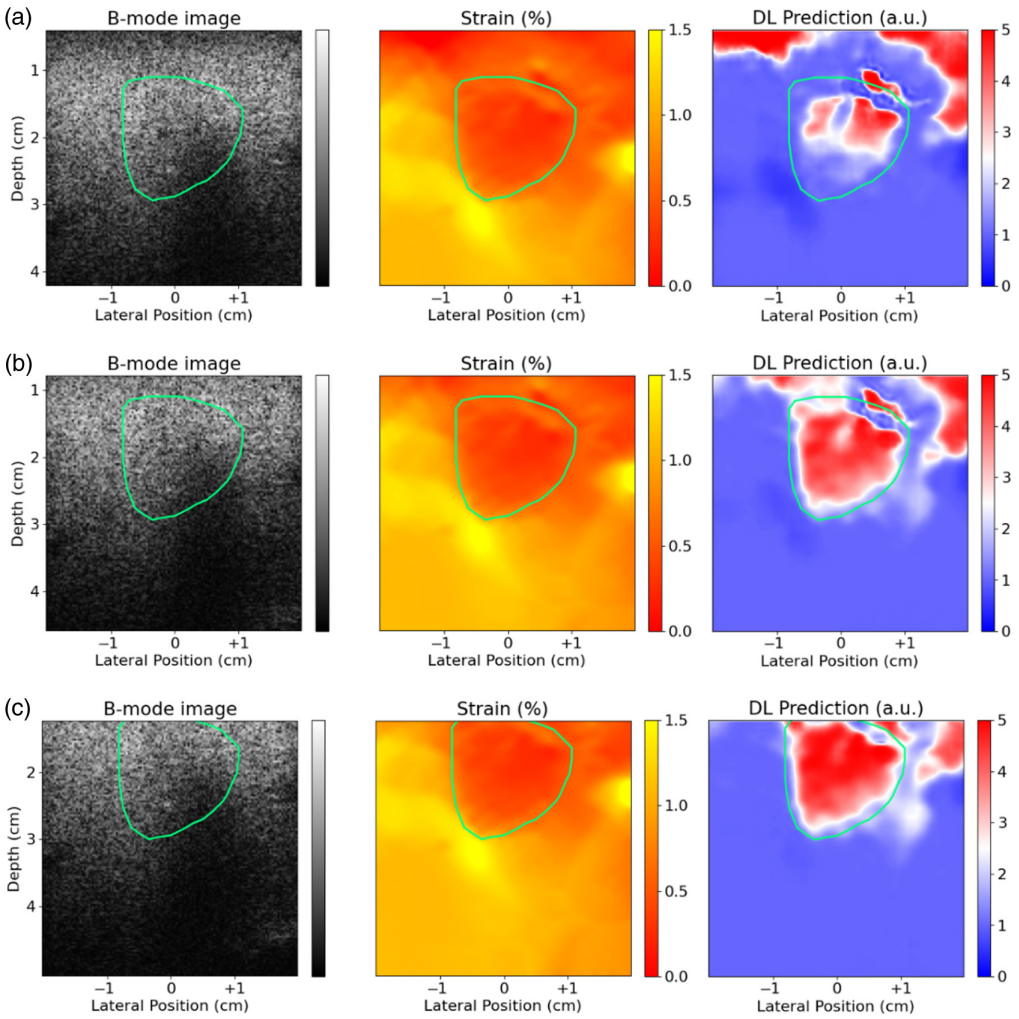

Fig. 13 Predictions for patient 1 for (a) first, (b) second, and (c) third ROIs.

4.3 Clinical Data

The clinical RF data were sourced from a publicly available dataset.[36] We applied our model to two patients from this dataset. The displacement fields were obtained using the image registration algorithm as described above. Upon applying our deep learning model to this, we observed notable results that merit discussion.

We showcase two patient's results with our model. Figures 13 and 14 demonstrate the prediction for the two patients respectively. The shadow observed in B-modes of Fig. 15 is caused by the thermal lesion. Here, it is important to note that the size and position of the lesions cannot be accurately determined from B-mode images alone,[36] but our algorithm is capable of identifying them based on their stiffness relative to that of the background tissue.

For patient 1, the modulus ratios were 3.60, 3.25, and 3.51 for the modulus images in Figs. 13(a)–13(c), with corresponding CNR values of 4.85, 4.57, and 5.05. For patient 2, the modulus ratios were 4.49, 3.33, and 3.15 for Figs. 14(a)–14(c), with CNR values of 7.14, 4.34, and 2.25, respectively. The output from our model also exhibited shapes that closely resembled those seen in their CT scans[36] which shows that our model can effectively capture the complex anatomical structures present in clinical scenarios.

5 Discussion

In this study, we generated a comprehensive simulated dataset designed to train a DL model and improve its ability to accommodate the experimental variability of the input data. This was achieved by simulating multiple domain sizes with inclusions at varying locations, from which

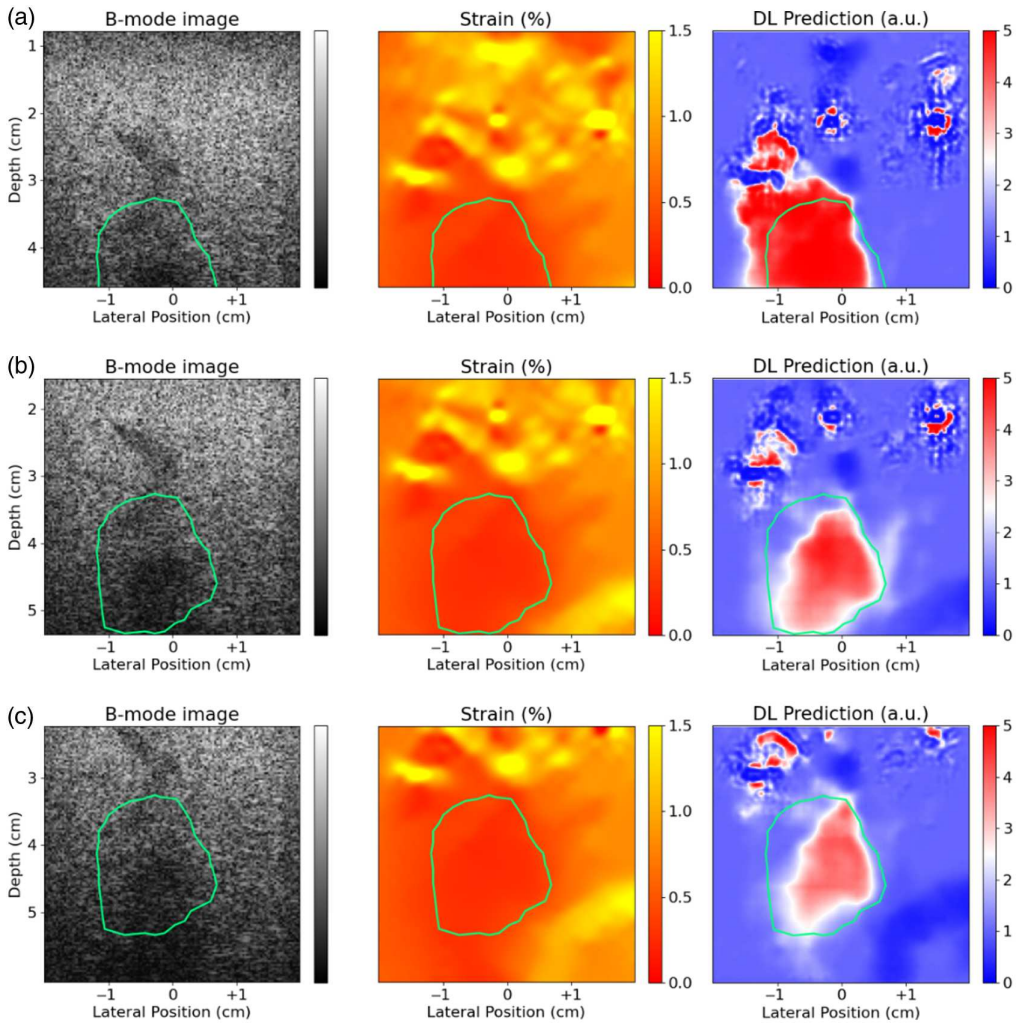

Fig. 14 Predictions for patient 2 for (a) first, (b) second, and (c) third ROIs.

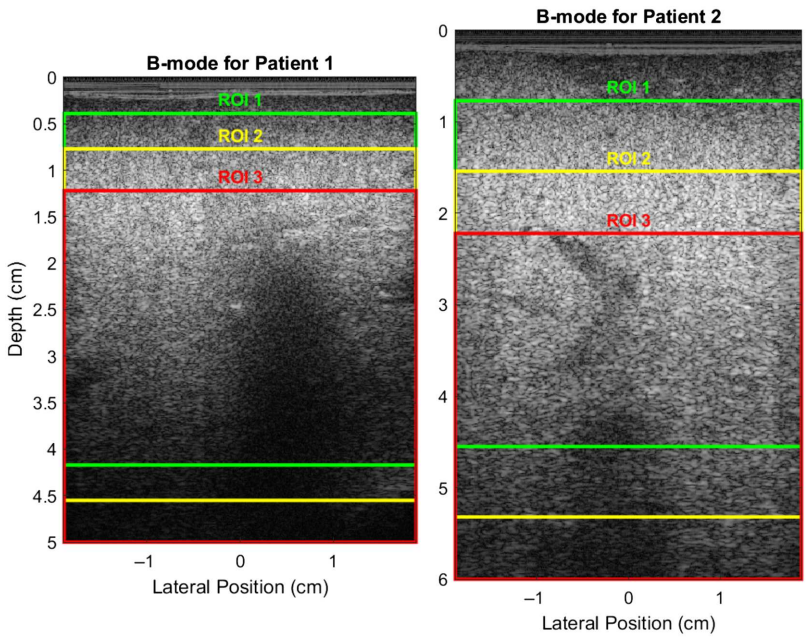

Fig. 15 B-mode of patients 1 and 2.

ROIs were extracted for model training, and introducing diverse boundary conditions, within those expected experimentally. The choice to split the ROIs was motivated by the need to manage computational resources and optimize the performance of our deep learning model. As the deep learning model requires, before even training, a specific choice of image dimension and with USE, the size, and aspect ratio of the pixel have a direct effect on the input and output data due to the direct connection to the mechanics, it made sense to begin with a ROI of fixed size (in mm) that can be moved within an image from the US to reconstruct our modulus images. Enlargement of the domain space allowed the incorporation of variability in boundary conditions and ROI. A notable benefit of this approach is evident when the inclusions are located at domain edges; the current model shows superior performance in these scenarios compared with the previous DL model, as detailed in Section 4.1.2. This is particularly significant for accurate modulus reconstruction at the edges within the simulated data. We also conducted a noise study by adding increasing amounts of Gaussian noise to our simulated test data. However, in real-world scenarios, noise can be more complex and varied than Gaussian noise. Ultrasound imaging systems also introduce system-dependent artifacts that may not be fully captured by a Gaussian noise model. To address this, we conducted experiments using data acquired using a real ultrasound system from a tissue-mimicking phantom. These experiments allowed us to evaluate the model under realistic noise conditions, providing insight into its performance when subjected to system-specific noise and artifacts. Despite the increased complexity of noise, such as those emulated by the phantom experiments, the algorithm maintains reasonable performance and reliability.

Our phantom study corroborates the accuracy of the deep-learning model, initially trained on simulated data, in handling experimental data. As shown in Table 3, the model's learned features from simulations prove transferable, enabling effective adaptation to phantom data. However, some residual noise was observed in modulus distribution predictions, possibly stemming from noisy input images or displacement estimations output by the registration algorithm. From ROI analysis as shown in Figs. 10 and 11, we can see that our DL model can handle edge cases. Further ROI analysis reveals that ROI 2 delivers optimal performance, indicating the model's robustness, particularly in reconstructing inclusions further away from the edges. Interestingly, ROI 1 consistently yields noisier predictions, which we hypothesize due to the proximity of the ultrasound transducer contact point to the ROI. This may be due to the overall low strain that exists at the surface of the transducer due to the low slip contact surface between the skin and the rigid transducer. Furthermore, we ranked the CNR across 40 samples for four different ROIs. The ranks are assigned in descending order of CNR quality, where rank 4 corresponds to the highest CNR value, indicating the best performance, whereas rank 1 represents the lowest CNR value, denoting the poorest performance. Based on the bar graph showing the average CNR ranks as shown in Fig. 16, we see that ROI 2 has the highest average rank, as it ranks fourth most of the time.

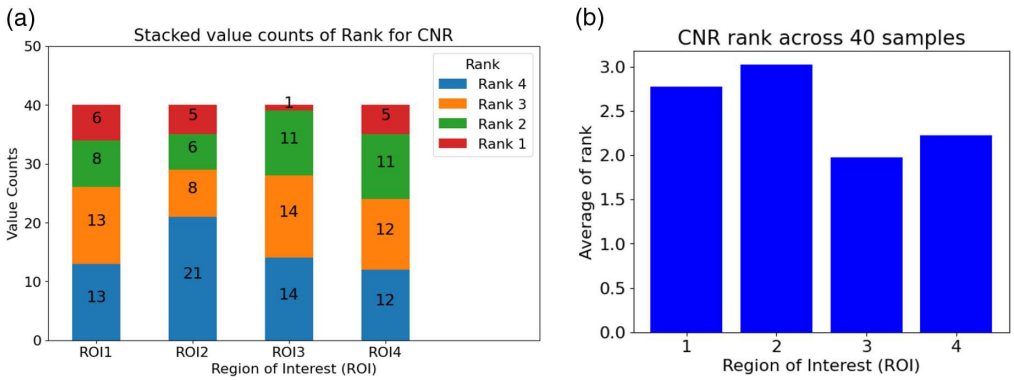

Fig. 16 Phantom ROI ranking. (a) ROI ranking frequency that provides a detailed enumeration of the rank for every ROI. Here, the ranks are assigned in descending order of CNR quality. Rank 4 corresponds to the highest CNR value, indicating the best performance, whereas rank 1 represents the lowest CNR value, denoting the poorest performance. (b) Average ranks of each ROI across all the samples.

We also present the application of our deep learning model to clinical data, a critical step toward translating our research into practical clinical utility. We can see more noise in the clinical data than in the simulated and phantom data. In general, the application of our deep learning model to clinical data has yielded promising results, demonstrating its potential for use in medical diagnostics. The similarities in shape with CT scans and the comparable contrast values underline the ability of the model to replicate and potentially augment traditional imaging techniques. These findings pave the way for further research and development in this field, with the ultimate goal of enhancing diagnostic precision and patient care in clinical settings. In practice, B-mode imaging does not always clearly show the lesions. There may be cases where the lesions appear faint, partially visible, or difficult to delineate entirely using the B-mode alone. For example, as shown in our clinical results, lesions are difficult to identify solely from images in B-mode. Using our method, we can detect regions of higher stiffness, supporting the utility of our approach in real-world clinical scenarios, especially even when the lesions are only partially visible or difficult to delineate using traditional B-mode images.

We recognize that the choice to use a regularized image registration technique, as opposed to a more computationally efficient displacement estimation technique, may seem counter to our long-term goal of real-time deployment. However, we hypothesize that using a more efficient algorithm would come with a higher degree of noise in the measurements and that this noise would likely need to be modeled in our training dataset. Thus, the CNN model would likely need to be trained in modeling specific US systems or system settings and thus to be less robust to deployment on different machines. We believe that this work is the first step toward understanding that trade-off.

For future work, our aim is to apply our model to training with both components of the displacement field. We also plan to test model accuracy with different displacement estimation techniques and compare the deep learning approach with traditional reconstruction techniques. In addition, we aim to improve our simulation to yield more realistic data for training that mimics the experimental US data better and perform noise studies for sensitivity.

One of the limitations of this work is that the training, validation, and test data are modeled as 2D plane strain. In addition, the displacement data are generated directly, rather than measured, and are generally noise-free, apparent from discretization error. In future models, we intend to investigate the feasibility of using 2D slices of geometrically three-dimensional models as our data and use the displacement data to simulate captured ultrasound images data, with added noise. Thus, the displacement fields used to train our NN can be measured with the same methods as our phantom and clinical data presented here. However, the computational complexity and cost of generating such data are much greater and thus left to future work. Another limitation is that the ROI is constrained to a square dimension here. In future work, we also aim to explore methods for processing larger ROIs while maintaining computational efficiency, such as utilizing patch-based approaches or scaling the model to handle larger images without compromising performance. We also plan on performing a broader hyperparameter tuning for our DL model. Finally, we hope to obtain a larger clinical dataset and perform a more rigorous feasibility study *in vivo*.

6 Conclusion and Future Work

We conducted a detailed study of how we can implement a deep learning framework to reconstruct the elastic modulus distribution from ultrasound-measured displacement fields. Our work includes a data simulation to train the DL model, along with controlled phantom studies and preliminary clinical results. With this, we demonstrate the model's capacity to generalize across variations and learn the relationship between displacement fields and modulus distribution. The model is also able to capture the variations in ROIs across the phantom data under different experimental conditions. This highlights that provided the simulated data is sufficiently diverse and representative of a wide variability, an algorithm trained on only simulated data would generalize well to real-world experimentally collected data. Finally, the clinical results offer initial insights into the applicability of our approach on data collected in a clinical setting. However, a larger patient dataset is necessary to draw definitive conclusions about the clinical efficacy of our approach, which will be the subject of our future work.

Disclosures

No conflicts of interest, financial, or otherwise.

Code and Data Availability

The data and the code used can be found on the GitHub repository here.

Acknowledgments

C. A. Linte and R. A. Simon have been supported by the National Institute of General Medical Sciences of the National Institutes of Health (Award No. R35GM128877), the Office of Advanced Cyber-infrastructure of the National Science Foundation (Award No. 1808530), and the Chemical, Bioengineering, Environmental and Transport Systems (CBET) of the National Science Foundation (Award No. 2245152).

References

1. J.-L. Gennisson et al., "Ultrasound elastography: principles and techniques," *Diagn. Interv. Imaging* **94**(5), 487–495 (2013).
2. R. M. Sigrist et al., "Ultrasound elastography: review of techniques and clinical applications," *Theranostics* **7**(5), 1303 (2017).
3. A. Goddi, M. Bonardi, and S. Alessi, "Breast elastography: a literature review," *J. Ultrasound* **15**(3), 192–198 (2012).
4. N. Frulio and H. Trillaud, "Ultrasound elastography in liver," *Diagn. Interv. Imaging* **94**(5), 515–534 (2013).
5. R. G. Barr et al., "Elastography assessment of liver fibrosis: society of radiologists in ultrasound consensus conference statement," *Radiology* **276**(3), 845–861 (2015).
6. E. Konofagou, "Cardiovascular elastography," in *Tissue Elasticity Imaging*, pp. 67–107, Elsevier (2020).
7. Y. Hong et al., "Real-time ultrasound elastography in the differential diagnosis of benign and malignant thyroid nodules," *J. Ultrasound Med.* **28**(7), 861–867 (2009).
8. J. Ophir et al., "Elastography: imaging the elastic properties of soft tissues with ultrasound," *J. Med. Ultrason.* **29**, 155–171 (2002).
9. J. Luo, K. Ying, and J. Bai, "Elasticity reconstruction for ultrasound elastography using a radial compression: an inverse approach," *Ultrasonics* **44**, e195–e198 (2006).
10. M. Bonnet and A. Constantinescu, "Inverse problems in elasticity," *Inverse Probl.* **21**(2), R1 (2005).
11. M. M. Doyley, "Model-based elastography: a survey of approaches to the inverse elasticity problem," *Phys. Med. Biol.* **57**(3), R35 (2012).
12. K. Raghavan and A. Yagle, "Forward and inverse problems in elasticity imaging of soft tissues," *IEEE Trans. Nucl. Sci.* **41**(4), 1639–1648 (1994).
13. P. Barbone et al., "Adjoint-weighted variational formulation for the direct solution of inverse problems of general linear elasticity with full interior data," *Int. J. Numer. Methods Eng.* **81**, 1713–1736 (2010).
14. N. Mohammadi, M. M. Doyley, and M. Cetin, "A statistical framework for model-based inverse problems in ultrasound elastography," in *54th Asilomar Conf. Signals, Systems, and Computers*, pp. 1395–1399 (2020).
15. M. Honarvar, R. Rohling, and S. E. Salcudean, "A comparison of direct and iterative finite element inversion techniques in dynamic elastography," *Phys. Med. Biol.* **61**(8), 3026 (2016).
16. H. Li et al., "Deep learning in ultrasound elastography imaging: a review," *Med. Phys.* **49**, 5993–6018 (2022).
17. M. G. Kibria and H. Rivaz, "Gluenet: ultrasound elastography using convolutional neural network," *Lect. Notes Comput. Sci.* **11042**, 21–28 (2018).
18. B. Peng, Y. Xian, and J. Jiang, "A convolution neural network-based speckle tracking method for ultrasound elastography," in *IEEE Int. Ultrason. Symp. (IUS)*, IEEE, pp. 206–212 (2018).
19. A. K. Tehrani and H. Rivaz, "Displacement estimation in ultrasound elastography using pyramidal convolutional neural network," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **67**(12), 2629–2639 (2020).
20. A. K. Tehrani and H. Rivaz, "MPWC-Net++: evolution of optical flow pyramidal convolutional neural network for ultrasound elastography," *Proc. SPIE* **11602**, 14–23 (2021).
21. S. Wu et al., "Direct reconstruction of ultrasound elastography using an end-to-end deep neural network," *Lect. Notes Comput. Sci.* **11070**, 374–382 (2018).
22. Z. Gao et al., "Learning the implicit strain reconstruction in ultrasound elastography using privileged information," *Med. Image Anal.* **58**, 101534 (2019).
23. D. Patel et al., "Circumventing the solution of inverse problems in mechanics through deep learning: application to elasticity imaging," *Comput. Methods Appl. Mech. Eng.* **353**, 448–466 (2019).

24. S. Cepeda et al., "Comparison of intraoperative ultrasound b-mode and strain elastography for the differentiation of glioblastomas from solitary brain metastases: an automated deep learning approach for image analysis," *Front. Oncol.* **10**, 590756 (2021).
25. L. He et al., "An application of super-resolution generative adversary networks for quasi-static ultrasound strain elastography: a feasibility study," *IEEE Access* **8**, 65769–65779 (2020).
26. C.-T. Chen and G. X. Gu, "Learning hidden elasticity with deep neural networks," *Proc. Natl. Acad. Sci. U. S. A.* **118**(31), e2102721118 (2021).
27. J. E. Warner et al., "Inverse estimation of elastic modulus using physics-informed generative adversarial networks," arXiv:2006.05791 (2020).
28. N. Mohammadi, M. M. Doyley, and M. Cetin, "Regularization by adversarial learning for ultrasound elasticity imaging," in *29th Eur. Signal Process. Conf. (EUSIPCO)*, IEEE, pp. 611–615 (2021).
29. N. H. Gokhale, "Solving an elastic inverse problem using convolutional neural networks," arXiv:2109.07859 (2021).
30. B. Ni and H. Gao, "A deep learning approach to the inverse problem of modulus identification in elasticity," *MRS Bull.* **46**, 19–25 (2020).
31. X. Zhang et al., "Displacement-based reconstruction of elasticity distribution with deep neural network," in *IEEE Int. Ultrason. Symp. (IUS)*, IEEE, pp. 1–5 (2022).
32. U. R. Tuladhar et al., "A deep learning framework to estimate elastic modulus from ultrasound measured displacement fields," *Proc. SPIE* **12470**, 124–130 (2023).
33. "C++ create random shaped 'blob' objects," StackOverflow, 2019, https://gamedev.stackexchange.com/questions/62613/need-ideas-for-an-algorithm-to-draw-irregular-blotchy-shapes (accessed 15 September 2015).
34. M. S. Richards and M. M. Doyley, "Non-rigid image registration-based strain estimator for intravascular ultrasound elastography," *Ultrasound Med. Biol.* **39**(3), 515–533 (2013).
35. M. S. Richards, Quantitative Three-Dimensional Elasticity Imaging, PhD thesis, Boston University (2007).
36. H. Rivaz et al., "Real-time regularized ultrasound elastography," *IEEE Trans. Med. Imaging* **30**(4), 928–945 (2010).
37. R. L. Chimenti et al., "Ultrasound strain mapping of Achilles tendon compressive strain patterns during dorsiflexion," *J. Biomech.* **49**(1), 39–44 (2016).
38. O. A. Babaniyi, A. A. Oberai, and P. E. Barbone, "Recovering vector displacement estimates in quasistatic elastography using sparse relaxation of the momentum equation," *Inverse Probl. Sci. Eng.* **25**(3), 326–362 (2017).
39. M. S. Richards, P. E. Barbone, and A. A. Oberai, "Quantitative three-dimensional elasticity imaging from quasi-static deformation: a phantom study," *Phys. Med. Biol.* **54**(3), 757 (2009).
40. O. Ronneberger, P. Fischer, and T. Brox, "U-Net: convolutional networks for biomedical image segmentation," *Lect. Notes Comput. Sci.* **9351**, 234–241 (2015).
41. P. Isola et al., "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 1125–1134 (2017).
42. S. L. Bangare et al., "Reviewing Otsu's method for image thresholding," *Int. J. Appl. Eng. Res.* **10**(9), 21777–21783 (2015).
43. R. I. of Technology, "Research computing services," (2019).
44. N. Anwar and F. Ahmed, "Structural cross-sections," (2017).

Utsav Ratna Tuladhar is an electrical and computer engineering PhD student at Rochester Institute of Technology, New York, United States. He completed his bachelor's degree in computer engineering from Tribhuvan University, Kathmandu, Nepal. His research interests include deep learning for medical applications, ultrasound elasticity imaging, and solving inverse problems.

Richard A. Simon, PhD, is a research scientist in the Department of Biomedical Engineering at Rochester Institute of Technology (Rochester, New York, United States). He holds a BSChem degree, with minors in mathematics and physics, and a doctorate in inorganic chemistry from Massachusetts Institute of Technology (Boston, Massachusetts, United States). Formerly, he was a senior research scientist at Carestream Health (Rochester, New York, United States) and Eastman Kodak (Rochester, New York, United States) and focused on various imaging-related projects.

Cristian A. Linte, PhD, is a professor in the Department of Biomedical Engineering and Chester F. Carlson Center for Imaging Science at Rochester Institute of Technology (Rochester, New York, United States). He holds a BAS degree in mechanical engineering, an MSE in biomedical

engineering, and a doctorate in biomedical engineering. His research focuses on the development of biomedical image computing, modeling, and visualization techniques for computer-integrated diagnosis and therapy.

Michael S. Richards, PhD, is an assistant professor of biomedical engineering at the Rochester Institute of Technology and an adjunct assistant professor in the Department of Surgery at the University of Rochester Medical Center. He holds a BS degree and a doctorate in biomedical engineering. His research focuses on studying disease-related changes to the biomechanics of soft tissues using ultrasound and other clinical imaging modalities.