

Real-time nonlinear inversion of magnetic resonance elastography with operator learning

Juampablo E. Heras Rivera, MSc*,¹ Caitlin M. Neher, MSc*,¹ and Mehmet Kurt, PhD¹

¹University of Washington, Seattle, WA, United States

Purpose: To develop and evaluate an operator learning framework for nonlinear inversion (NLI) of brain magnetic resonance elastography (MRE) data, which enables real-time inversion of elastograms with comparable spatial accuracy to NLI.

Materials and Methods: In this retrospective study, 3D MRE data from 61 individuals (mean age, 37.4 years; 34 female) were used for development of the framework. A predictive deep operator learning framework (oNLI) was trained using 10-fold cross-validation, with the complex curl of the measured displacement field as inputs and NLI-derived reference elastograms as outputs. A structural prior mechanism, analogous to Soft Prior Regularization in the MRE literature, was incorporated to improve spatial accuracy. Subject-level evaluation metrics included Pearson’s correlation coefficient, absolute relative error, and structural similarity index measure between predicted and reference elastograms across brain regions of different sizes to understand accuracy. Statistical analyses included paired t-tests comparing the proposed oNLI variants to the convolutional neural network baselines.

Results: In comparison to convolutional architectures, oNLI showed superior accuracy when used to reconstruct both the storage and loss moduli (μ' and μ'' , respectively) from complex curl fields. Whole brain absolute percent error was $8.4\% \pm 0.5\%$ (μ') and $10.0\% \pm 0.7\%$ (μ'') for oNLI and $15.8\% \pm 0.8\%$ (μ') and $26.1\% \pm 1.1\%$ (μ'') for CNNs. Additionally, oNLI outperformed convolutional architectures as per Pearson’s correlation coefficient r in the whole brain and across all subregions for both the storage modulus and loss modulus ($p < .05$). In the thalamus, oNLI achieved a mean μ' correlation of $r = 0.92$ versus $r = 0.87$ for CNNs. In the hippocampus, oNLI achieved $r = 0.84$ versus $r = 0.65$. Across white matter, oNLI achieved $r = 0.93$ versus $r = 0.78$, and in the cortex, $r = 0.92$ versus $r = 0.75$. Finally, oNLI predictions of μ'' significantly outperformed CNNs in the whole brain ($r = 0.96$ versus 0.82), cerebral cortex ($r = 0.95$ versus 0.79), white matter ($r = 0.96$ versus 0.81), thalamus ($r = 0.95$ versus 0.88), and hippocampus ($r = 0.91$ versus 0.68).

Conclusion: The oNLI framework enables real-time MRE inversion (30,000x speedup), outperforming CNN-based approaches and maintaining the fine-grained spatial accuracy achievable with NLI in the brain.

1. Introduction

Magnetic resonance elastography (MRE) [1] is a quantitative method for noninvasively obtaining the mechanical properties of sub-superficial organs of interest, such as the liver, brain, and kidneys [2–4]. In the liver, MRE is the most accurate noninvasive technique for diagnosing and staging fibrosis [5], and has become a routine clinical scan. In the brain, a growing body of evidence suggests that viscoelasticity derived from MRE is sensitive to aging and disease [6, 7] and regional properties correlate with age [8], cognitive decline [9], ApoE status [10], and perfusion [11].

MRE involves three components: the mechanical actuation of the tissue of interest, a phase-contrast MR pulse sequence to encode three-directional displacements during actuation, and post-processing using an inversion algorithm to recover mechanical properties of the tissue from the displacement data. Frequently reported properties are shear stiffness and damping ratio, which provide information on the structural integrity and viscosity of the tissue of interest. Numerous inversion algorithms have been developed to relate tissue displacement to mechanical properties, each with different material model assump-

tions, computational costs, and final property estimates.

When a linear and locally homogeneous material is assumed, the complex shear modulus, μ , is defined by the Helmholtz equation:

$$\mu = \frac{-\rho\omega^2\mathbf{u}}{\nabla^2\mathbf{u}} \quad (1)$$

where $\mathbf{u}$ is the complex harmonic displacement vector (or its curl), ρ is the density, and ω is the frequency of mechanical actuation. Inversion approaches based on this assumption used in practice include direct inversion [12], local frequency estimation [13], and algebraic inversion of the differential equation [14]. Applying this material model simplifies the inversion and allows for fast computations on the order of seconds, but leads to artifacts near boundaries and inaccuracies around small features. This limits the applicability of these approaches, as interfaces and small anatomical regions like the sub-cortical gray matter cannot be distinctly resolved.

By removing the local homogeneity assumption, the complex shear modulus is defined by the heterogeneous form of the time-harmonic Navier’s equation, which describes the evolution of displacement fields in a viscoelastic medium as:

$$\nabla \cdot (\mu(\nabla\mathbf{u} + \nabla\mathbf{u}^T)) + \nabla(\lambda\nabla \cdot \mathbf{u}) = -\rho\omega^2\mathbf{u}, \quad (2)$$

where $\mathbf{u}$ is the complex harmonic displacement vector, μ is the complex viscoelastic shear modulus, λ is the second Lamé parameter, ρ is the density, ω is the frequency of mechanical actuation. The most common inversion approach used for this material model is nonlinear inversion (NLI), a finite-element-based iterative method that estimates viscoelastic parameters by minimizing the difference between measured and simulated displacement fields [15, 16]. NLI can characterize the sub-cortical material property changes associated with learning [17], cognition [18, 19], and epileptic foci [20]. Additionally, heterogeneity allows for the accurate representation of stiffness at interfaces [21]. While NLI is better suited for organs with high tissue complexity, and makes fewer assumptions, it is several orders of magnitude slower than direct inversion. The high computational cost of NLI, which requires more than four hours of runtime using 32 CPU cores for a standard $132 \times 132 \times 60$ matrix size, makes it difficult to implement clinically when real-time results are desired.

Here, we frame inversion of (2) as done in NLI as a neural operator learning [22] problem. We term this method as operator NLI (oNLI), where the goal is to learn a mapping between MRE displacement fields and complex shear stiffness in the infinite-dimensional function space setting. Unlike finite-element based methods which require an iterative optimization from scratch for each subject, oNLI learns a general operator that maps any subject’s MRE displacement/curl fields to their respective shear moduli in one forward pass.

Abbreviations

MRE: magnetic resonance elastography, NLI: nonlinear inversion, oNLI: operator learning framework for NLI, SSIM: structural similarity index measure, CNN: convolutional neural network, SPADE: spatially-adaptive normalization.

Summary

The operator learning framework for NLI (oNLI) enables inversion in fractions of a second, while maintaining the fine-grained spatial accuracy achievable with NLI in the brain.

Key Points

- oNLI, a deep operator learning framework, was developed for brain MRE nonlinear inversion using curl fields as inputs and NLI elastograms as reference outputs. Additionally, a mechanism for incorporating structural priors was introduced.
- Across 61 subjects and 10-fold cross-validation, oNLI achieved higher Pearson’s correlation with ground truth NLI than convolutional neural networks in all major brain subregions (e.g., μ' , white matter: $r = 0.93$ vs. 0.78 ; μ'' , hippocampus: $r = 0.91$ vs. 0.68).
- oNLI reconstructed elastograms $30,000\times$ faster than conventional NLI (fractions of a second vs. multiple hours) while maintaining fine-grained spatial accuracy.

Keywords

Magnetic Resonance Elastography, Operator Learning, Brain, Nonlinear Inversion, Deep Learning, SPADE, Convolutional Neural Networks.

By framing NLI as an operator learning problem, oNLI provides a resolution-invariant mapping, overcoming the issues of convolutional neural network (CNN) approaches, which fail to generalize beyond one imaging resolution. Although the upfront training cost is high, once deployed oNLI enables real-time nonlinear inversion (< 1 second on one GPU, < 30 seconds on one CPU), representing a **30,000x** speedup over NLI. We train instances of oNLI using an openly available 3D brain MRE

dataset, and demonstrate that elastograms produced by oNLI show excellent visual and quantitative agreement with ground-truth elastograms when evaluated on held-out subjects. Finally, we introduce SPADE-oNLI, a variant for the encoding of spatial priors relevant to the inversion process, similar to the Soft Prior Regularization approach introduced in [23].

2. Methods

2.1. Study Design

This work is a retrospective study using previously collected data. The primary goal is to design a framework for training deep operators to perform nonlinear inversion of magnetic resonance elastography data in real-time.

2.2. Data

Dataset Description

This study used a publicly available dataset from the Brain Biomechanics Imaging Resources in the Neuroimaging Tools and Resources Collaboratory (NITRC) for training and validation. The specific data used was collected at the University of Delaware, where 61 subjects were scanned on a Siemens Prisma 3T scanner with a 64 channel head and neck coil. Each subject dataset includes an anatomical T1 image at 0.8 mm isotropic resolution and a 3D MRE multiband, multishot spiral sequence acquisition at 30, 50, and 70 Hz at 1.5 mm isotropic resolution. The MRE displacements were generated in the anterior-posterior direction using the Resoundant pneumatic actuation system with a soft pillow driver. High-resolution mechanical property maps were obtained using the NLI algorithm [16].

The subjects used for training and validation are between 14 to 75 years old, with 34F/27M and an age distribution described by the histogram in Figure 2. These 61 subjects have a mean age of 37.43 ± 20.46 years.

For this retrospective study, data were included based on public availability. To our knowledge, all subjects were without neurological conditions at the time the data was collected.

Data Format

For each subject and actuation frequency, complex-valued displacement fields from MRE acquisitions and their derived curl fields, along with complex shear modulus maps (storage and loss moduli) estimated

via NLI were used. Anatomical T1-weighted scans were coregistered to the MRE reference space to enable voxel-wise correspondence to MRE fields.

Inputs to oNLI

For each spatial location $\mathbf{x}$ on the imaging domain $\Omega \subset \mathbb{R}^3$, the complex-valued displacement field $\mathbf{u}(\mathbf{x}) \in \mathbb{C}^3$ was obtained from the temporal Fast Fourier Transform (FFT) [26] of the measured displacement $\mathbf{r}(\mathbf{x}, t)$, $t \in \mathbb{R}$, at the actuation frequency ω :

$$\mathbf{u}(\mathbf{x}) = \frac{1}{T} \int_0^T \mathbf{r}(\mathbf{x}, t) e^{-i\omega t} dt. \quad (3)$$

The real and imaginary components of $\mathbf{u}(\mathbf{x})$ were each stored as arrays of shape $(160 \times 160 \times 80 \times 3)$, where the first three dimensions index the spatial grid and the last dimension corresponds to the three displacement directions (x, y, z) . From this displacement field, the curl $(\nabla \times \mathbf{u}(\mathbf{x}))$ is obtained, which aids inversion by suppressing longitudinal waves and enhancing the shear wave components produced by mode conversion. For all experiments, the real and imaginary components of the curl fields were concatenated along the fourth dimension, together with an additional channel containing the actuation frequency (in Hz) divided by 100. This yielded arrays of shape $(160 \times 160 \times 80 \times 7)$, which were used as inputs to the oNLI framework.

Outputs from oNLI

The outputs from oNLI are complex shear modulus fields $\mu(\mathbf{x}) = \mu'(\mathbf{x}) + i\mu''(\mathbf{x}) \in \mathbb{C}^3$, where μ' and μ'' correspond to the storage and loss moduli, respectively. These fields are represented as four-dimensional arrays of shape $(160 \times 160 \times 80 \times 2)$, where the first three dimensions index the spatial grid and the fourth dimension indexes the real and imaginary components of the complex shear modulus.

Missing Data

Five NITRC subject datasets were excluded completely due to missing MRE data. Additionally, one subject was missing 30 Hz MRE data and four subjects were missing 70 Hz MRE data. These subjects with partially missing data were not excluded; instead, their available data were used for training and validation.




Figure 1: **a)** oNLI overview: Complex 3D displacement maps, $\mathbf{u}(\mathbf{x}) = (u_x, u_y, u_z) \in \mathbb{C}^3$ for $\mathbf{x}$ on the imaging domain $\mathbf{x} \in \Omega \subset \mathbb{R}^3$, are input to the neural operator, their curl is computed, lifted to a higher dimension by linear mapping $\mathcal{P}$, processed by T Operator layers, and projected to the complex shear modulus, $\mu(\mathbf{x}) \in \mathbb{C}$, by the linear map $\mathcal{Q}$. **b)** Diagrams of operator layers used in this work to instantiate oNLI. Top: Diagram of Fourier Layer introduced in [24], where the operator kernel is parametrized in the Fourier domain. Bottom: Diagram of the SPADE layer used in the SPADE-oNLI variant. Anatomical T1 scans are processed using SynthSeg [25] to obtain anatomical segmentations, which are passed through convolutional layers to extract spatially varying statistics, γ and β . These statistics are then used to modulate the affine parameters of instance normalization.




Figure 2: Histogram showing the age distribution of the subjects used for training and validation, including male (blue) and female (orange) partitions for each bin.

2.3. Ground Truth

The ground truth material property maps were computed using the nonlinear inversion algorithm developed in [16]. NLI iteratively solves for the shear modulus as described by the heterogeneous form of the time-harmonic Navier’s equation (Equation (2)).

2.4. Data Partitions

To evaluate model architectures and extract performance statistics, 10-fold cross-validation was performed. The dataset contains scans from 61 subjects, each at multiple frequencies, so the folds were divided on a per-subject basis to avoid data leakage between training and validation sets. The dataset was randomly split into 10 folds such that each fold contained

54-55 subjects in the training set and 5-6 in the validation set.

2.5. Model

Neural Operators for MRE Inversion

We introduce **oNLI** (operator nonlinear inversion), a neural operator framework for MRE. Given complex curl fields $\nabla \times \mathbf{u}(\mathbf{x}) \in \mathbb{C}^3$ defined on the imaging domain $\Omega \subset \mathbb{R}^3$, the goal is to recover the spatially varying shear modulus $\mu(\mathbf{x}) \in \mathbb{C}^3$.

Formally, let $\mathcal{U}$ denote the function space of complex curl fields $\nabla \times \mathbf{u} : \Omega \rightarrow \mathbb{C}^3$ and $\mathcal{Y}$ the function space of complex modulus fields $\mu : \Omega \rightarrow \mathbb{C}^3$. The inversion task is expressed as learning the operator $\mathcal{G} : \mathcal{U} \rightarrow \mathcal{Y}$, defined by

$$\mathcal{G}(\nabla \times \mathbf{u})(\mathbf{x}) = \mu(\mathbf{x}), \quad \mathbf{x} \in \Omega, \quad (4)$$

where $\mathcal{G}$ is resolution-invariant and mesh-free. Unlike finite-dimensional networks that depend on voxel discretization, neural operators approximate mappings between function spaces, allowing a single set of learned parameters to generalize across resolutions and subjects. This is highly beneficial for MRE, where voxel sizes vary across scanners and protocols.

In practice, $\mathcal{G}$ is implemented by composing learnable kernel operators $\{\mathcal{K}_i\}_{i=1}^T$ with pointwise nonlinearities, where T denotes the number of kernel layers. Each kernel operator has the form

$$(\mathcal{K}_i v)(x) = \int_{\Omega} \kappa_i(x, y) v(y) dy, \quad (5)$$

such that

$$\mathcal{G} = \mathcal{Q} \circ \sigma \circ \mathcal{K}_T \circ \sigma \circ \dots \circ \sigma \circ \mathcal{K}_1 \circ \mathcal{P}, \quad (6)$$

where σ denotes a pointwise nonlinearity (e.g., ReLU), v denotes an intermediate feature field, and $\mathcal{P}$ and $\mathcal{Q}$ are learned linear maps that lift the input curl $\nabla \times \mathbf{u}$ to a higher-dimensional feature space and project intermediate representations back to the complex modulus field, respectively. A diagram of the oNLI framework is shown in Figure 1.

Fourier Neural Operator

To instantiate oNLI, we adopt the Fourier neural operator (FNO). FNOs are efficient approximators as they parametrize the operator kernel in the Fourier domain. Given input v , a Fourier layer computes

$$(\mathcal{K}(v))(x) = \mathcal{F}^{-1}(R \cdot \mathcal{F}(v))(x), \quad (7)$$

where $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote the Fourier transform and its inverse, respectively, and R is a linear transform on the lower Fourier modes. By combining global spectral convolution with local nonlinearities, FNOs approximate highly nonlinear PDE-driven operators with quasi-linear complexity [27]. Importantly, since the operator is learned in continuous space, the same parameters apply across discretizations, allowing for super-resolution and fast inference (millisecond-scale per volume) without retraining or finetuning.

Incorporating anatomical priors

T1-weighted MRI acquired alongside MRE provides high-resolution anatomical structure information. We incorporate these priors into oNLI using whole-brain segmentations generated by **SynthSeg** [25], which is substantially faster than **recon-all** from FreeSurfer [28] while maintaining acceptable accuracy. Segmentations from **SynthSeg** include over 33 anatomical regions. We group these into six broader categories: cortical gray matter, white matter, and subcortical gray matter, which have statistically distinct material properties [29]; as well as cerebrospinal fluid (CSF), brainstem/cerebellum, and background, which are typically noisy in MRE data and thus treated as separate regions. Segmentation masks are represented as a one-hot encoded vector, resulting in a 6-channel binary segmentation mask of shape $\{0, 1\}^{6 \times H \times W \times D}$, where H , W , and D denote the spatial dimensions. Figure 3 shows a representative sagittal slice of a T1-weighted scan and the corresponding anatomical segmentation mask.




Figure 3: T1-weighted MRI (left) and corresponding 6-region segmentation mask (right) generated using SynthSeg. Regions include: background (black), cortical gray matter (red), white matter (blue), subcortical gray matter (green), brainstem/cerebellum (purple), and cerebrospinal fluid (CSF; orange).

To integrate these anatomical priors into the oNLI framework, we adopt the SPADE [30] approach, a spatially adaptive normalization method originally developed for semantic image synthesis. SPADE modulates the affine parameters of instance normalization using spatially varying statistics derived from the segmentation map. In SPADE-oNLI, the multi-channel segmentation mask $S \in \{0, 1\}^{6 \times H \times W \times D}$ is processed by a small convolutional network to predict spatially adaptive scale $\gamma(S)$ and bias $\beta(S) \in \mathbb{R}^{C \times H \times W \times D}$ maps. Given an intermediate feature map $F \in \mathbb{R}^{C \times H \times W \times D}$, the modulated output is:

$$\text{SPADE}(F, S) = \gamma(S) \cdot \text{IN}(F) + \beta(S)$$

where $\text{IN}(F)$ is the instance-normalized version of F before affine transformation. A diagram illustrating a SPADE-oNLI layer is shown in Figure 1b). This conditioning is effective in preserving high-frequency features, counteracting the low-frequency bias of neural operators observed in [31] and imposing structural priors. We observe that this layer-wise conditioning strategy preserves the conditioning throughout the model inference much better than passing the conditioning as an additional input. Unlike the original SPADE paper, which recommends 3x3 convolutions tied to a fixed grid resolution, we use 1x1 convolutions, so that conditioning is independent of spatial discretization, preserving oNLI’s resolution invariant property. For the SPADE-oNLI variant considered in the experiments, the operator kernel was parametrized in the Fourier domain as in the Fourier Neural Operator.

2.6. U-Net

As a baseline comparison, the 3D U-Net [32] architecture was used. U-Net has been widely adopted for medical image processing tasks due to its strong inductive bias for spatially localized features. The model consists of an encoder-decoder structure with

skip connections that preserve spatial detail across resolutions. Unlike oNLI, which learns an operator between function spaces, U-Net learns a finite-dimensional function mapping defined on discretized grids. As a result, the model is tied to a fixed discretization, and inversion of inputs with different resolutions requires resampling to match the resolution of the training data.

Model architecture

Three model architectures were considered for evaluation of the proposed method, namely: U-Net (baseline), oNLI, and SPADE-oNLI. All models were configured with seven input channels: 6 channels for the real and imaginary curl components in three orthogonal directions, and 1 channel for frequency conditioning. All models had two output channels, corresponding to the real and imaginary parts of the complex shear modulus. For all experiments, the number of model parameters was fixed to between 84–85 million to enable a fair comparison between methods while ensuring compatibility with smaller GPUs and maintaining practical inference times (sub-second on GPU, < 30 seconds on CPU).

The baseline U-Net model was configured with 4 encoder levels with channels (84, 168, 336, 672), 3 downsampling steps (stride 2), and residual blocks with instance normalization.

The oNLI and SPADE-oNLI models were both configured with an FNO parametrization of the operator kernel. However, we emphasize that the oNLI formulation is not limited to Fourier kernels; in principle, oNLI can be implemented with alternative operator kernels, such as wavelets. In both models, each layer used 5 stacked layers, each with 20 Fourier modes per spatial dimension and width 23. The SPADE-oNLI variant used the 3D SPADE block at each layer, with input channels equal to the number of classes in the one-hot anatomical segmentation, a single hidden convolutional layer of 32 channels, and a final convolutional layer with 23 channels to match the FNO layer width. For all convolutional layers, instance normalization and kernel size 1 were used to preserve resolution invariance.

2.7. Training

Data normalization

Inputs and targets were normalized independently for curl and stiffness fields using per-channel unit Gaussian standardization. Specifically, we applied the

UnitGaussianNormalizer implementation from the **neuraloperators** library [27, 33], with the mean and standard deviation estimated from the training set. After forward passes, model predictions were denormalized using the inverse of the preprocessing step, to preserve stable training while still recovering the full variability in magnitude scales across subjects and frequencies.

Training Details

Following the hyperparameter choices in [34], all models were trained using the Adam optimizer with an initial learning rate of 10^{-3} and a weight decay of 10^{-4} . The cosine annealing learning rate scheduler with $T_{\max} = \text{epochs} \times \lfloor N_{\text{train}} / \text{batch size} \rfloor$ was used. All models were trained for 50 epochs with a batch size of 1.

The relative L_2 loss was used as an optimization metric, as specified in [27]. For predictions $\mu^{oNLI} \in \mathbb{R}^{N \times M}$ and ground-truth fields $\mu^{\text{NLI}} \in \mathbb{R}^{N \times M}$, where N is the batch size and M is the number of voxels per sample, the loss is defined as

$$\mathcal{L}_2^{\text{data}}(\mu^{oNLI}, \mu^{\text{NLI}}) = \frac{1}{N} \sum_{i=1}^N \frac{\|\mu_i^{oNLI} - \mu_i^{\text{NLI}}\|_2}{\|\mu_i^{\text{NLI}}\|_2}. \quad (8)$$

Here, $\|\cdot\|_2$ denotes the L_2 -norm. This formulation normalizes the prediction error by the target norm to preserve scale invariance across subjects and frequencies.

Compute resources and software used

Each model was trained using 8 CPU cores and a single NVIDIA A40 GPU on the University of Washington’s Hyak supercomputing cluster. Model development was done using PyTorch [35], and data analysis with Python.

2.8. Evaluation

In this study, the predictions from three candidate models: U-Net (baseline), oNLI, and SPADE-oNLI were compared to the NLI ground truth reference. Model evaluation was performed using the held-out datasets from each fold in the 10-fold cross-validation. Predictions from all the held-out subjects across all folds were concatenated to form a pooled validation set, which was used to calculate group-level evaluation metrics. Analyses were conducted on individual brain subregions including: the cerebral cortex,

Table 1: Model performance by parameter (μ' and μ'') and region (whole brain, cerebral cortex, white matter, thalamus, and hippocampus) across pooled validation subjects. Pearson’s correlation coefficient (r), absolute percent error (APE; %), and structural similarity index measure (SSIM) are reported. SPADE-oNLI outperformed U-Net in the whole brain and the selected regions except for the SSIM of μ' for the whole brain, in which U-Net performed the best.

Model	Storage Modulus (μ')			Loss Modulus (μ'')		
	Pearson’s r ($\uparrow$)	APE (%) ($\downarrow$)	SSIM ($\uparrow$)	Pearson’s r ($\uparrow$)	APE (%) ($\downarrow$)	SSIM ($\uparrow$)
Whole Brain						
U-Net	0.782 ± 0.023	15.8 ± 0.8	0.720 ± 0.043	0.822 ± 0.020	26.1 ± 1.1	0.520 ± 0.041
oNLI	0.907 ± 0.012	9.4 ± 0.6	0.670 ± 0.041	0.944 ± 0.008	10.7 ± 0.7	0.595 ± 0.049
SPADE-oNLI	0.934 ± 0.008	8.4 ± 0.5	0.679 ± 0.040	0.956 ± 0.006	10.0 ± 0.7	0.598 ± 0.051
Cerebral Cortex						
U-Net	0.752 ± 0.026	16.5 ± 0.8	0.731 ± 0.066	0.786 ± 0.023	25.5 ± 1.1	0.510 ± 0.072
oNLI	0.883 ± 0.015	10.6 ± 0.6	0.726 ± 0.047	0.933 ± 0.010	10.8 ± 0.7	0.680 ± 0.050
SPADE-oNLI	0.915 ± 0.012	9.6 ± 0.6	0.735 ± 0.041	0.946 ± 0.008	10.8 ± 0.7	0.685 ± 0.046
White Matter						
U-Net	0.775 ± 0.024	15.2 ± 0.8	0.741 ± 0.067	0.806 ± 0.022	28.6 ± 1.2	0.560 ± 0.042
oNLI	0.909 ± 0.011	9.1 ± 0.5	0.755 ± 0.052	0.949 ± 0.007	11.1 ± 0.7	0.711 ± 0.055
SPADE-oNLI	0.932 ± 0.010	7.6 ± 0.6	0.760 ± 0.051	0.955 ± 0.007	10.1 ± 0.8	0.715 ± 0.054
Thalamus						
U-Net	0.871 ± 0.015	17.8 ± 0.9	0.438 ± 0.205	0.881 ± 0.015	31.1 ± 1.8	0.203 ± 0.137
oNLI	0.898 ± 0.012	16.1 ± 1.1	0.408 ± 0.166	0.932 ± 0.009	17.3 ± 1.3	0.390 ± 0.144
SPADE-oNLI	0.915 ± 0.011	14.1 ± 1.3	0.442 ± 0.166	0.946 ± 0.007	15.6 ± 1.3	0.403 ± 0.144
Hippocampus						
U-Net	0.650 ± 0.037	16.4 ± 0.9	0.734 ± 0.105	0.676 ± 0.036	25.5 ± 1.3	0.629 ± 0.094
oNLI	0.833 ± 0.020	11.3 ± 0.6	0.732 ± 0.081	0.893 ± 0.019	14.3 ± 1.0	0.721 ± 0.085
SPADE-oNLI	0.841 ± 0.018	10.6 ± 0.7	0.754 ± 0.086	0.910 ± 0.013	11.8 ± 0.8	0.734 ± 0.083

white matter, thalamus, and hippocampus. These regions are commonly studied using NLI and vary significantly in size; here they are used to provide an estimate of model performance on small and large regions. The group-level evaluation metrics reported include: the correlation coefficient between predicted and ground truth values, the absolute relative error, and the structural similarity index (SSIM) between the predicted and corresponding ground truth shear modulus elastograms. All group-level metrics are reported in Table 1, and analyzed in the Section 3. In addition to the group-level evaluation metrics, we report the fold-level statistics, including the mean, standard deviation, and 95% confidence interval of the validation relative L_2 loss across folds.

3. Results

3.1. Population Demographics

A total of 61 subjects (34 female and 27 male; mean age, 37.43 years, standard deviation 20.46 years) were

included in the model development and evaluation dataset. Among these participants, all were without neurological conditions at the time the data was collected. In the following subsections, fold-wise and group-wise metrics are reported. For group-wise metrics, the predictions for the held-out datasets from all trained models in the 10-fold cross validation are compiled into a larger validation dataset, hereafter referred to as the “pooled validation set”.

3.2. Cross-Validation Performance

The validation relative L_2 loss (Equation (8)) for each of the folds ($n = 10$) of all models is shown in Table 2. The average loss and standard deviation for all folds are 0.338 ± 0.008 , 0.298 ± 0.014 , and 0.291 ± 0.012 for U-Net, oNLI, and SPADE-oNLI, respectively. Paired t-tests across folds showed that both oNLI ($p \approx 3.7 \times 10^{-5}$) and SPADE-oNLI ($p \approx 9.0 \times 10^{-7}$) achieved significantly lower validation loss than U-Net, while the difference between

Table 2: Mean validation relative L_2 loss (Equation (8)) per fold, with across-fold mean, standard deviation (STD), and 95% confidence interval (CI).

Model	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5	Fold 6	Fold 7	Fold 8	Fold 9	Fold 10	Mean	STD	95% CI
U-Net	0.337	0.332	0.325	0.330	0.336	0.351	0.347	0.344	0.329	0.344	0.338	0.008	[0.332, 0.344]
oNLI	0.280	0.321	0.291	0.305	0.282	0.285	0.314	0.297	0.289	0.316	0.298	0.014	[0.289, 0.307]
SPADE-oNLI	0.270	0.293	0.291	0.289	0.288	0.289	0.298	0.293	0.281	0.319	0.291	0.012	[0.283, 0.299]




Figure 4: Ground truth vs. predicted mean storage and loss moduli (μ' and μ'' , respectively) across validation subjects ($n = 56$) at three actuation frequencies (30, 50, and 70 Hz). Pearson's r with respect to a linear regression fit is reported for each comparison. Across the regions of interest (cerebral cortex, white matter, thalamus, and hippocampus), SPADE-oNLI performs the best with respect to Pearson's r .

oNLI and SPADE-oNLI was not statistically significant ($p \approx 0.075$). The oNLI variants both outperform the U-Net model across all individual folds, indicating a robust improvement. The low average standard deviation indicates similar performance across folds for the full dataset with no outlier folds. Finally, the 95% confidence intervals for both oNLI variants are lower than that of the U-Net, indicating with high certainty that performance has improved.

3.3. Prediction vs. Ground Truth Across Frequencies

Figure 4 shows a comparison between ground truth and predicted values for the storage and loss moduli

(μ' and μ'' , respectively) across validation subjects in the pooled validation set at three frequencies. Each subplot shows scatter plots of predicted vs. ground truth values for individual regions of interest (cerebral cortex, white matter, thalamus, and hippocampus), with a blue dashed regression line and corresponding r value to indicate the correlation of the model's predictions with its trend line. Additionally, the $y = x$ reference ideal line is shown for visual comparison, and the correlation coefficients between the model predictions and the ground truth are shown in Table 1.




Figure 5: Absolute percent error of the storage and loss moduli (μ' and μ'' , respectively) as evaluated against ground truth mean values for the cerebral cortex, white matter, thalamus and hippocampus. oNLI shows a significant error reduction from U-Net across all regions and both moduli except for the storage modulus in the thalamus, and SPADE-oNLI shows a significant error reduction from U-Net across all cases. ns: $p \geq 0.5$; *: $p < 0.05$; **: $p < 0.01$; ***: $p < 0.001$; ****: $p < 0.0001$.




Figure 6: **a)** Comparison of ground truth and predicted storage modulus (μ'), **b)** loss modulus (μ''), **c)** magnitude of the complex shear modulus ($|\mu| = \sqrt{(\mu')^2 + (\mu'')^2}$) and **d)** damping ratio ($\xi = \mu''/2\mu'$) for one representative validation subject at 50Hz. Within each **a-d)**, from left to right, columns show the respective parameter derived from NLI (ground truth), U-Net prediction, oNLI prediction, and SPADE-oNLI prediction, respectively. From top to bottom: coronal, axial, and sagittal slices are shown for each **a-d)**.

For the U-Net variants, we observe that at 30 Hz, U-Net is able to produce estimates which are near the $y = x$ reference line; however, at 50 and 70 Hz, U-Net systematically underpredicts the mean stiffness values across all subregions, even though frequency was included as a conditioning input. Furthermore, U-Net struggles to accurately model the loss modulus, and predicts near constant mean values across frequencies, indicating that this model architecture is unable to capture the monotonically increasing stiffness-frequency relationship.

For the baseline oNLI variant and the SPADE-oNLI variant we observe that the model predictions closely approach the $y = x$ reference line across all frequencies. Both variants accurately capture the increas-

ing stiffness-frequency relationship in both the storage and loss moduli. Quantitatively, values of r for SPADE-oNLI outperform U-Net in all regions analyzed for μ' (whole brain: 0.93 versus 0.78; cerebral cortex: 0.92 versus 0.75; white matter: 0.93 versus 0.78; thalamus: 0.92 versus 0.87; and hippocampus: 0.84 versus 0.65) and μ'' (whole brain: 0.96 versus 0.82; cerebral cortex: 0.95 versus 0.79; white matter: 0.96 versus 0.81; thalamus: 0.95 versus 0.88; and hippocampus: 0.91 versus 0.68).

3.4. Regional Relative Error Analysis

Figure 5 shows a comparison of the absolute percent relative errors (APE) for all subjects in the pooled validation set across brain subregions for all considered models. This measure provides an estimate for the

error that can be expected from using each model to predict the mean value of the storage and loss moduli in each of the subregions. All APE values are presented in Table 1 as well. Paired t-tests were used to perform pairwise comparisons between models for each region, and significance is reported in Figure 5. Across the cerebral cortex, white matter, and hippocampus regions, for both storage and loss modulus, we observe that the oNLI variants significantly outperform the U-Net variants with $p < 0.0001$. For the storage modulus in the thalamus, we see that the baseline oNLI variant is not significantly better than the U-Net, but SPADE-oNLI is significantly better than both ($p < 0.05$ vs. baseline oNLI and $p < 0.01$ vs. U-Net). Generally, we observe that SPADE-oNLI significantly outperforms U-Net on both small and large brain regions.

3.5. Structural Similarity Index Measure (SSIM)

The SSIM metric was used to further quantify each model’s full-brain and regional performance. Whole brain analysis showed that U-Net performed the best for predictions of μ' with a SSIM of 0.72. However, SPADE-oNLI achieved the best SSIM for the rest of the μ' predictions (cerebral cortex, white matter, thalamus and hippocampus) and all of the μ'' predictions, including the whole brain. Interestingly, the U-Net predictions of the storage modulus were comparable to SPADE-oNLI across all subregions, but the loss modulus predictions were significantly worse. Since SSIM is a metric which combines luminance, contrast and structure, it provides a more holistic understanding of the spatial distribution of values. From the elastograms in Figure 6, we see that U-Net correctly predicts the relative storage modulus in the ventricles, white matter, and gray matter. However, it predicts the same spatial distribution at a different scale for the loss modulus, failing to learn a pattern distinct from the storage modulus. This effect is highlighted when the damping ratio is plotted in Figure 6d), where the voxel-wise calculation $\mu''/2\mu'$ results in an almost completely constant and incorrect elastogram. This occurs because CNNs use the same learned filters across channels, limiting their ability to predict significantly different moduli at different channels.

4. Discussion

This study introduces oNLI, a data-driven operator approach to perform MRE inversion in fractions of a

second in heterogeneous materials with an accuracy comparable to nonlinear inversion. We develop this approach using 3D brain MRE data and show that operator learning is significantly more data-efficient during training than convolutional network-based approaches such as U-Net for this application, achieving superior accuracy at multiple actuation frequencies when trained on a dataset of only 61 subjects. Furthermore, we introduce SPADE-oNLI, which allows for the incorporation of spatial priors during inference, further improving the predictions made by oNLI. The proposed inversion requires minimal computational overhead once trained, and results in a 30,000x speedup over traditional NLI methods. All presented approaches require no additional scans or scan time, which aid in seamless integration for future iterations.

By developing a framework which directly predicts heterogeneous material properties, this work represents progress towards real-time high-fidelity MRE in a clinical diagnostic setting and broadens the diagnostic capability of rapid inversion beyond the liver to more complex organs such as the brain. This reduced inversion time has the potential to broaden the use of MRE in both research and clinical settings, accelerating the development of NLI-based biomarkers. When trained on a dataset with various resolutions, we hypothesize that the oNLI approach will further outperform U-Net variants, given the continuous nature of the representation learned.

Our study had several limitations, which should be addressed in future works. The proposed model was trained and evaluated using retrospective data from a single site, so further analysis is required to validate the results in prospective multicenter studies. Additionally, training and evaluation using multi-resolution data should be done to further utilize the capabilities of resolution-invariant operator learning models. Furthermore, the methods presented were only evaluated with a healthy cohort, and should be further developed to ensure robust performance in characterizing pathological tissue properties.

In conclusion, this study proposes a data-driven operator learning approach for real-time inversion of Navier’s equation with the heterogeneous material assumptions, as done with NLI. We introduce a method to introduce anatomical priors into the model training by using SPADE layers, a concept typically used for semantic image generation. By leveraging anatomical priors, the proposed approach is able to closely

match the ground truth NLI predictions with reduced variance and improved accuracy in comparison to U-Net and the FNO-based oNLI. The presented models robustly and significantly outperform U-Net with respect to Pearson’s r , the APE, and the SSIM as evidenced by 10-fold cross-validation. Future work will primarily focus on a) collecting a larger and more diverse dataset for training and evaluation of the proposed method on multiple sites and resolutions, b) evaluating the proposed approach on clinical data, and c) incorporating physics-informed simulation data to supplement the limited training dataset size.

5. Author Contributions

J.E.H.R.: conceptualization, data curation, formal analysis, investigation, methodology, software, validation, writing—original draft, writing—review and editing; C.M.N.: conceptualization, data curation, formal analysis, investigation, methodology, software, validation, writing—original draft, writing—review and editing; M.K.: funding acquisition, supervision, project administration, resources, writing—review.

6. Acknowledgments

The authors thank Dr. Curtis Johnson’s group for providing the NITRC MRE dataset and Dr. Matt McGarry for sharing the NLI implementation.

7. Funding Statement

The work of Juampablo Heras Rivera was partially supported by the U.S. Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0024386. The work of Caitlin Neher was partially supported by the National Science Foundation under Award Number NSF CMMI 1953323.

References

- [1] R. Muthupillai, D. J. Lomas, P. J. Rossman, J. F. Greenleaf, A. Manduca, and R. L. Ehman. Magnetic resonance elastography by direct visualization of propagating acoustic strain waves. *Science*, 269(5232):1854–1857, 1995.
- [2] Raja Muthupillai, David J. Lomas, Philip J. Rossman, James F. Greenleaf, Armando Manduca, and Richard L. Ehman. Magnetic resonance elastography by direct visualization of propagating acoustic strain waves. *Science*, 269(5232):1854–1857, 1995.
- [3] Armando Manduca, T. E. Oliphant, M. A. Dresner, J. L. Mahowald, S. A. Kruse, E. Amromin, J. P. Felmlee, J. F. Greenleaf, and R. L. Ehman. Magnetic resonance elastography: Non-invasive mapping of tissue elasticity. *Medical Image Analysis*, 5(4):237–254, December 2001.
- [4] Olivier Rouvière, Rémi Souchon, Gaële Pagnoux, Jean-Michel Ménager, and Jean-Yves Chapelon. Magnetic resonance elastography of the kidneys: feasibility and reproducibility in young healthy adults. *Journal of Magnetic Resonance Imaging*, 34(4):880–886, oct 2011.
- [5] Siddharth Singh, Sudhakar K. Venkatesh, Zhen Wang, Frank H. Miller, Utaroh Motosugi, Yoshito Kamiya, and et al. Magnetic resonance elastography for staging liver fibrosis in nonalcoholic fatty liver disease: A pooled individual participant data analysis. *Hepatology*, 62(3):748–756, 2015.
- [6] Han Lv, Mehmet Kurt, Nan Zeng, Erdal Ozkaya, Fabien Marcuz, Lyndia Wu, Kaveh Laksari, David B. Camarillo, Kim Butts Pauly, Zhenchang Wang, and Max Wintermark. MR elastography frequency-dependent and independent parameters demonstrate accelerated decrease of brain stiffness in elder subjects. *European Radiology*, 30(12):6614–6623, 2020.
- [7] Yuan Feng, Matthew C. Murphy, Emi Hojo, Fei Li, and Neil Roberts. Magnetic resonance elastography in the study of neurodegenerative diseases. *Journal of Magnetic Resonance Imaging*, 59(1):82–96, 2024.
- [8] Lucy V. Hiscox, Hillary Schwarb, Matthew D. J. McGarry, and Curtis L. Johnson. Aging brain mechanics: Progress and promise of magnetic resonance elastography. *NeuroImage*, 232:117889, 2021.
- [9] KowsalyaDevi Pavuluri, John Huston, Richard L. Ehman, Armando Manduca, Prashanthi Vemuri, Clifford R. Jack, Matthew L. Senjem, and Matthew C. Murphy. Brain mechanical properties predict longitudinal cognitive change in aging and alzheimer’s disease. *Neurobiology of Aging*, 147:203–212, 2025.
- [10] Lucy V Hiscox, Bradley Karat, Louisa E Wood, Robert C Davis, Anne Corbett, Claudia Metzler-Baddeley, Curtis L Johnson, and Derek K Jones. MR elastography reveals lower hippocampal stiffness in middle-aged apoe ϵ 4 carriers without cognitive impairment. 2025.
- [11] Caitlin Maria Neher, Em Triolo, Fargol Rezayaraghi, Oleksandr Khagai, Priti Balchandani, Matthew McGarry, and Mehmet Kurt. Perfusion-mechanics coupling of the hippocampus. *Interface Focus*, 15(1):20240051, April 2025. Published April 4, 2025.
- [12] A. J. Romano, J. J. Shirron, and J. A. Bucaro. On the noninvasive determination of material parameters from a knowledge of elastic displacements theory and numerical simulation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 45(3):751–759, May 1998.
- [13] A. Manduca, R. Muthupillai, P. J. Rossman, J. F. Greenleaf, and R. L. Ehman. Local wavelength estimation for magnetic resonance elastography. In *Proceedings of 3rd IEEE International Conference on Image Processing*, volume 3, pages 527–530, Lausanne, Switzerland, 1996.
- [14] Travis E Oliphant, Armando Manduca, Richard L Ehman, and James F Greenleaf. Complex-valued stiffness reconstruction for magnetic resonance elastography by algebraic inversion of the differential equation. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 45(2):299–310, 2001.

- [15] E. E. W. Van Houten, M. I. Miga, J. B. Weaver, F. E. Kennedy, and K. D. Paulsen. Three-dimensional subzone-based reconstruction algorithm for MR elastography. *Magnetic Resonance in Medicine*, 45:827–837, 2001.
- [16] Matthew D McGarry, Elijah E Van Houten, Carter L Johnson, John G Georgiadis, Bradley P Sutton, John B Weaver, and Keith D Paulsen. Multiresolution MR elastography using nonlinear inversion. *Medical Physics*, 39(10):6388–6396, October 2012.
- [17] H. J. Heseltun, A. T. Anderson, C. L. Johnson, N. J. Cohen, B. P. Sutton, and H. Schwarb. Putamen stiffness declines with age and is associated with implicit sequence learning outcomes. *Brain Sciences*, 15(9):947, 2025.
- [18] Peyton L. Delgorio, Lucy V. Hiscox, Grace McIlvain, Mary K. Kramer, Alexa M. Diano, Kyra E. Twohy, Alexis A. Merritt, Matthew D. J. McGarry, Hillary Schwarb, Ana M. Daugherty, James M. Ellison, Alyssa M. Lanzi, Matthew L. Cohen, Christopher R. Martens, and Curtis L. Johnson. Hippocampal subfield viscoelasticity in amnesic mild cognitive impairment evaluated with MR elastography. *NeuroImage: Clinical*, 37:103327, 2023.
- [19] Peyton L. Delgorio, Lucy V. Hiscox, Ana M. Daugherty, Faria Sanjana, Grace McIlvain, Ryan T. Pohlig, Matthew D. J. McGarry, Christopher R. Martens, Hillary Schwarb, and Curtis L. Johnson. Structure-function dissociations of human hippocampal subfield stiffness and memory performance. *Journal of Neuroscience*, 42(42):7957–7968, 2022.
- [20] Graham R. Huesmann, Hillary Schwarb, Daniel R. Smith, Ryan T. Pohlig, Aaron T. Anderson, Matthew D. J. McGarry, Keith D. Paulsen, Tracey Mencion Wszalek, Bradley P. Sutton, and Curtis L. Johnson. Hippocampal stiffness in mesial temporal lobe epilepsy measured with MR elastography: Preliminary comparison with healthy participants. *NeuroImage: Clinical*, 27:102313, 2020.
- [21] Chengyuan Wu, Mahdi Alizadeh, Mary K. Kramer, Matthew B. Kroen, Robert Ziechmann, Feroze B. Mohamed, Qianhong Wu, and Curtis L. Johnson. Deep brain stimulation electrode deviations are associated with brain stiffness interfaces measured by magnetic resonance elastography. *Operative Neurosurgery*, page ons.0000000000001523, 2025. Published online February 20, 2025.
- [22] Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89):1–97, 2023.
- [23] Matthew McGarry, Courtney L. Johnson, Bradley P. Sutton, Elijah E. Van Houten, John G. Georgiadis, John B. Weaver, and Keith D. Paulsen. Including spatial information in nonlinear inversion MR elastography using soft prior regularization. *IEEE Transactions on Medical Imaging*, 32(10):1901–1909, Oct 2013. Epub 2013 Jun 17.
- [24] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations, 2021.
- [25] Benjamin Billot, Douglas N. Greve, Oula Puonti, Axel Thielscher, Koen Van Leemput, Bruce Fischl, Adrian V. Dalca, and Juan Eugenio Iglesias. SynthSeg: Segmentation of brain mri scans of any contrast and resolution without retraining. *Medical Image Analysis*, 86:102789, 2023.
- [26] James W Cooley and John W Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of computation*, 19(90):297–301, 1965.
- [27] Nikola B. Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces. *CoRR*, abs/2108.08481, 2021.
- [28] Bruce Fischl. Freesurfer. *NeuroImage*, 62(2):774–781, 2012.
- [29] Lucy V Hiscox, Matthew DJ McGarry, Hillary Schwarb, Elijah EW Van Houten, Ryan T Pohlig, Neil Roberts, Graham R Huesmann, Agnieszka Z Burzynska, Bradley P Sutton, Charles H Hillman, et al. Standard-space atlas of the viscoelastic properties of the human brain. *Human brain mapping*, 41(18):5282–5300, 2020.
- [30] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2337–2346, 2019.
- [31] Shaoxiang Qin, Fuyuan Lyu, Wenhui Peng, Dingyang Geng, Ju Wang, Xing Tang, Sylvie Leroyer, Naiping Gao, Xue Liu, and Liangzhu Leon Wang. Toward a better understanding of fourier neural operators from a spectral perspective. *arXiv preprint arXiv:2404.07200*, 2024.
- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *CoRR*, abs/1505.04597, 2015.
- [33] Jean Kossaifi, Nikola Kovachki, Zongyi Li, Davit Pitt, Miguel Liu-Schiaffini, Robert Joseph George, Boris Bonev, Kamyar Azizzadenesheli, Julius Berner, and Anima Anandkumar. A library for learning neural operators, 2024.
- [34] Wenhui Peng, Shaoxiang Qin, Senwen Yang, Jianchun Wang, Xue Liu, and Liangzhu (Leon) Wang. Fourier neural operator for real-time simulation of 3D dynamic urban microclimate. *Building and Environment*, 248:111063, 2024.
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.