Published in final edited form as:

Magn Reson Med. 2018 July ; 80(1): 351–360. doi:10.1002/mrm.27019.

Artificial neural networks for stiffness estimation in magnetic resonance elastography

Matthew C. Murphy¹, Armando Manduca^{1,2}, Joshua D. Trzasko¹, Kevin J. Glaser¹, John Huston III¹, and Richard L. Ehman¹

¹Department of Radiology, Mayo Clinic, Rochester, MN, USA

²Department of Physiology and Biomedical Engineering, Mayo Clinic, Rochester, MN, USA

Abstract

Purpose—To investigate the feasibility of using artificial neural networks to estimate stiffness from MR elastography (MRE) data.

Methods—Artificial neural networks were fit using model-based training patterns to estimate stiffness from images of displacement using a patch size of approximately 1 cm in each dimension. These neural network inversions (NNIs) were then evaluated in a set of simulation experiments designed to investigate the effects of wave interference and noise on NNI accuracy. NNI was also tested *in vivo*, comparing NNI results against currently used methods.

Results—In four simulation experiments, NNI performed as well or better than direct inversion (DI) for predicting the known stiffness of the data. Summary NNI results were also shown to be significantly correlated with DI results in the liver ($R^2=0.974$) and in the brain ($R^2=0.915$), and also correlated with established biological effects including fibrosis stage in the liver and age in the brain. Finally, repeatability error was lower in the brain using NNI compared to DI, and voxel-wise modeling using NNI stiffness maps detected larger effects than using DI maps with similar levels of smoothing.

Conclusion—Artificial neural networks represent a new approach to inversion of MRE data. Summary results from NNI and DI are highly correlated and both are capable of detecting biologically relevant signals. Preliminary evidence suggests that NNI stiffness estimates may be more resistant to noise than an algebraic DI approach. Taken together, these results merit future investigation into NNIs to improve the estimation of stiffness in small regions.

Keywords

MR elastography; inversion; stiffness; neural networks; shear modulus

Introduction

Magnetic resonance elastography (MRE) is an MRI-based technique for measuring tissue stiffness (1), akin to manual palpation, which has a long history in the practice of medicine.

It is a three-step process beginning with the introduction of shear waves into the tissue of interest by mechanical vibration. Next, the resulting tissue displacement is imaged by a phase-contrast MRI pulse sequence with motion encoding gradients that are synchronized to the motion. Finally, these displacement images are mathematically inverted to calculate tissue stiffness. MRE is growing in clinical impact, particularly in the liver, where it is used to noninvasively assess fibrosis with high accuracy (2–5). Furthermore, brain stiffness has shown sensitivity to a number of physiological processes including neurodegenerative disease (6–10), normal aging (11,12), and even behavioral performance (13,14). Still, significant motivation remains to improve the effective resolution and robustness of the technique to enable further investigation into the spatial relationship between tissue mechanical properties and pathophysiology, and to open up new applications in focal diseases.

Currently, the inversion algorithms used to compute stiffness from the displacement data represent one of the greatest limitations to the resolution of MRE. A number of methods exist for performing this calculation, each with its unique set of advantages and disadvantages. Popular algorithms in the field include, but are not limited to, variations of direct inversion (DI) (15–17), local frequency estimation (18–20), and non-linear inversion (21–23). Each model relies on underlying assumptions that will impact its performance in *in vivo* data, as well as a number of implementation choices (e.g., tuning parameters and filter settings) that affect the tradeoff between resolution and numerical stability. An algorithm that can provide stable estimates in the presence of noise while maintaining a small spatial filtering footprint could substantially improve the effective resolution of MRE, either by mitigating the need for smoothing or by allowing reliable estimation in smaller regions.

Given the limitations of currently used algorithms, this work investigates the feasibility of a new approach to MRE inversion using artificial neural networks (ANNs). In general, a machine learning approach to MRE inversion is appealing because it is data-driven, providing a straightforward framework for dealing with practical issues (e.g., tissue inhomogeneity, partial volume effects and noise), and predictions made from these models are effectively bounded by the training data, providing the stability necessary to measure stiffness in smaller regions. ANNs in particular are interesting for this application because we do not know *a priori* the optimal features to extract from the displacement data in order to estimate stiffness. ANNs with multiple hidden layers are capable of learning abstract representations of the inputs, which are then utilized by deeper layers to estimate the quantity of interest (i.e., end-to-end learning) (24). Therefore, this manuscript presents initial results for utilizing ANNs to estimate stiffness from displacement data, beginning with simulation experiments, and followed by applications in *in vivo* liver and brain data.

Theory

Interest in ANNs has rapidly accelerated in recent years as they have outperformed competing approaches in several difficult tasks including image classification and speech recognition (24). Here we present only a brief overview of ANNs to better understand this work. Comprehensive background on this technique is available elsewhere such as (25).

ANNs represent a class of function approximators inspired by neural networks found in the brain. These networks consist of a set of model neurons such as the one shown schematically in Fig. 1A. The essential properties of this model neuron are that it, first, computes the weighted sum of input from upstream neurons (x) as well as a bias term (b), and second, applies a nonlinear transfer function to this value, $h(x)$, and sends this output to downstream neurons. A hyperbolic tangent transfer function (also called an activation function) was used in this work, though several other candidate functions are also commonly used. While a number of ANN architectures exist, this work utilizes feedforward networks with fully-connected layers. An example feedforward ANN is shown in Figure 1B, which has 3 inputs (or features), 2 hidden layers with 5 units in each, and finally a single neuron in the output layer. In this type of ANN, each neuron receives input from every neuron in the previous layer, and sends its output to every neuron in the following layer. Each of these neurons operates as the model shown in Figure 1B. There are not strict criteria for deciding on the appropriate size of an ANN for a particular application. In general, more neurons will give the model increased capacity to fit complex, nonlinear functions at the cost of increased risk for overfitting. Overfitting can be avoided by providing additional training examples, or implementing techniques such as regularization or dropout (26).

Finally, the weights of the ANN must be learned through the process of training. In the case of supervised training, the weights are adjusted iteratively by the method of backpropagation using a training set consisting of features and known desired outputs (27). First, the weights are initialized typically with small, random values. Then the ANN is evaluated in the forward direction, and the error is calculated between the ANN prediction and the known target values provided by the training set. Next, the gradient of the error with respect to each weight is computed in the backward direction beginning with the output layer and working progressively upstream, each time using the previously computed gradients to calculate errors in the immediately upstream-layer according to the chain rule. With the error and gradients computed, any one of a number of nonlinear optimization algorithms can be used to update the weights for the next iteration. This process is repeated until a stopping criterion is reached. Due to the non-convexity of such problems, resulting from the use of nonlinear transfer functions ($h(x)$), only local optimality is guaranteed. However, this limitation often does not restrict the efficacy of such methods in practice. In this work, training was stopped when error was no longer reduced in a separate cross-validation set.

Methods

Simulation experiments

The feasibility of a neural network inversion (NNI) was first examined in four different sets of 2D simulated data with increasing complexity. All analyses were performed in MATLAB R2016a (Mathworks, Natick, MA). These experiments followed same overall design:

1. Generate 140,000 simulated data sets (100,000 for training, 20,000 for validation, and 20,000 for testing).
2. Extract relevant features from the displacement data.

3. Fit the feedforward neural network model. All models used in this work had 3 hidden layers of 24 units per layer with a hyperbolic tangent transfer function, and were fit using scaled conjugate gradient backpropagation (28). Network size was determined by gradually increasing model capacity until performance plateaued in simulations. Training was stopped when mean squared error was not improved in the validation set for 6 consecutive iterations.
4. Evaluate performance in test set. The neural network model as trained above was run on the test data sets. A linear model was fit between the estimated stiffness from the neural network and the true stiffness for each example in the test set, and this relationship was compared with the same result from DI. Since DI can produce large outliers that obscure this relationship, examples where DI estimated stiffness greater than 20 kPa were excluded for regression analysis of the DI results only.

The simulated data sets were 5×5 patches with 2-mm isotropic voxels and 4 phase offsets spaced evenly over one period of 60-Hz motion. The stiffness of each example was chosen randomly from a uniform distribution with a range of 0.1 to 10 kPa, where stiffness is defined as the product of density and the square of wave speed (20). For the simplest simulation, a single point-source was placed in a random location within a circular shell at distance from the center of the patch of 1–2× the stiffest possible wavelength. A sinusoidal wave was computed as a function of distance from that source to produce radially propagating waves with the appropriate wavelength (20), and no attenuation. The features given to the model for each observation, x_i , were the real and imaginary parts of the first temporal harmonic of the displacement data in this 5×5 patch (1 cm × 1 cm of displacement data), making a total of 50 features. The features of each example were linearly scaled such that the range of $\pm|x_i|$ was mapped to the range –1 to +1.

The remaining three experiments gradually added complexity to the simulations. The second experiment added multiple wave sources. Each training example had a random number of sources in the range of 1 to 10, each placed at a random location within the above circular shell. The final simulated displacements were then simply the superposition of the multiple sources divided by the number of sources. For the third experiment, we again considered single-source simulations where each training example was given a randomly chosen signal-to-noise ratio (SNR) taken from a uniform distribution in the range of 1 to 20. Pseudorandom noise with a Gaussian distribution was generated for each data set with zero mean and the standard deviation scaled to provide the prescribed SNR, and added to the displacement data (i.e., the simulated MRE phase difference signal before calculating the first temporal harmonic). Finally, the fourth experiment included both multiple wave sources and noise. In all cases, the same features and patch scaling were used as in the first experiment.

In vivo experiments

All human studies were performed after obtaining approval from our institutional review board and written informed consent. We considered two *in vivo* applications including the liver and the brain. In general, these experiments began by fitting a neural network using

simulated data, with the assumptions underlying those simulations chosen specific to each application. Models for use *in vivo* were fit with 1 million examples for training and 200,000 for validation. The model was then applied to *in vivo* data and the results were compared to standard MRE approaches. Details related to each application are below.

Liver experiment

The liver NNI was trained on 11×11 2D patches with 1-mm isotropic voxels and 4 phase offsets. Simulation parameters included stiffness in the range of 0.1 to 10 kPa for 60-Hz motion, number of wave sources in the range of 1 to 10 placed randomly in a circular shell (distance from center of the patch of 1–2× wavelength at 10 kPa), and SNR in the range of 1 to 20 (noise added to data before taking temporal first harmonic). The liver NNI utilized 242 features that include the real and imaginary parts of the first harmonic displacement data in an 11×11 patch (maintaining the roughly 1 cm × 1 cm footprint used in simulation experiments), with each set of features scaled to the range of –1 to +1 as above.

MRE exams from 86 subjects for this portion of the work were taken from a previous study (29). MRE data were acquired on a 1.5-T MR scanner (Signa, GE Medical Systems, Milwaukee, WI) with a modified gradient echo sequence with the following parameters: 60-Hz vibration, TR/TE = 50/20.2 ms, 30° flip angle, FOV in the range of 30 to 48 cm; 256×64 imaging matrix zero-padded to 256×256, 2× ASSET acceleration, 4 contiguous 10-mm thick axis slices, one 16.7-ms motion-encoding gradient, slice-direction motion encoding, and 4 phase offsets spaced evenly over one period of 60-Hz motion.

The data were first inverted by multiple-model direct inversion (MMDI), a variant of direct inversion commonly used for clinical liver MRE exams (30). MMDI is typically performed following 2D directional filtering to mitigate the effects of interfering waves (31). This inversion algorithm fits the displacement data with multiple polynomial models and chooses the best fitting model with an F-test. The R^2 values from this modeling serve as a measure of inversion confidence. MMDI also utilizes total least squares fitting when solving the Helmholtz equation (32). NNI stiffness maps were then computed by evaluating the NNI at every voxel of the liver exams using wave images after phase unwrapping, removal of background phase (performed by the MMDI algorithm by subtracting the estimated constant term at each voxel), and 2D directional filtering. For both MMDI and NNI, four directional filters were used with 90° spacing and cutoffs of 2 and FOV/2 waves per FOV. To account for the different resolutions between patients, all wave images were interpolated to 1-mm isotropic in-plane resolution using cubic spline interpolation prior to inversion by the neural network. In simulation results, we observed variations in the stiffness maps related to the phase of the underlying displacement data, likely stemming from the amplitude of the spatial derivative of the displacements. However, once the first harmonic is estimated, the phase can be set arbitrarily. Therefore, to reduce any potential bias in the NNI result related to the phase of the displacement data in a particular patch, NNI was run 32 times with the phase of the input data varied linearly over one cycle of motion and the 32 stiffness maps were averaged. This process is analogous to cycle spinning, which is commonly used in wavelet denoising to mitigate artifact (33). Both inversions were summarized by computing the mean stiffness in a region of interest (ROI) that was the intersection of an expert-drawn ROI and

the MMDI confidence map thresholded at an R^2 of 0.95. For comparison, the results were fit by a restricted cubic spline model with 3 knots with 95% confidence intervals computed by bootstrapping. To determine if NNI can detect the effects of fibrosis as has been previously established for MMDI, we also display summary data as a function of fibrosis stage and report the R^2 value of a linear model that predicts stiffness by fibrosis stage (treated as a categorical variable). For display only, NNI stiffness maps were smoothed by a circular median filter with a radius of 7 pixels to roughly match the apparent resolution of MMDI stiffness maps, while maps without smoothing are provided in the Supplementary Material.

Brain experiment

The brain NNI was trained on $5 \times 5 \times 5$ 3D patches with 3-mm isotropic voxels and 8 phase offsets. Simulation parameters included stiffness in the range of 0.1 to 5 kPa for 60-Hz motion, wave sources in the range of 1 to 10 placed randomly in a spherical shell (distance from center of patch of $1-2 \times$ wavelength at 5 kPa), and SNR in the range of 1 to 20 (noise added to data before taking temporal first harmonic). The brain NNI utilized 54 features that include the real and imaginary parts of the first harmonic displacement data in a $3 \times 3 \times 3$ patch taken from the center of the simulated patch, scaled as described above.

Data for this portion of the work were taken from previous studies (11,34). MRE data were acquired using a 3-T MR scanner (Signa, GE Medical Systems, Milwaukee, WI) with a modified spin-echo echo planar imaging sequence with the following parameters: 60-Hz vibration; TR/TE = 3600/62 ms; FOV = 24 cm; 72×72 imaging matrix zero-padded to 80×80 ; 48 contiguous 3-mm thick axis slices; one 18.2-ms motion-encoding gradient on each side of the refocusing pulse; x, y, and z motion encoding directions; and 8 phase offsets spaced evenly over one period of 60-Hz motion. A 3D T1-weighted image was also collected to segment the brain and co-register an atlas to the MRE data. This image was acquired with an IR-SPGR pulse sequence with the following parameters: sagittal orientation, superior-inferior frequency encoding direction, TR/TE = 6.3/2.8 ms, 11° flip angle, TI = 400 ms, FOV = 27 cm, 256×256 imaging matrix, BW = ± 31.25 kHz, $1.75 \times$ ASSET acceleration, and 200 slice locations with 1.2-mm spacing.

From the displacement data, we used edge-aware methods as previously described to compute the curl, and then smooth the resulting curl data with a $3 \times 3 \times 3$ quartic kernel without using data from outside the prescribed ROI (34). These curl images were then inverted using the algebraic DI approach (16), and stiffness was summarized as the median over the ROI after eroding the mask by one voxel. To avoid passing the inversion algorithms missing data from outside the brain, the curl data were effectively dilated by one voxel. This operation was performed by smoothing the curl data with a $3 \times 3 \times 3$ median filter (but excluding voxels outside the brain from the calculation), and then creating a composite curl image that was equal to the unsmoothed curl data inside the brain and equal to the smoothed estimate in voxels outside the brain. NNI stiffness maps were then computed by evaluating the trained neural network in every voxel in the ROI using the first temporal harmonic of the curl as the inputs to the trained ANN. Similarly to the liver NNI, the brain NNI was run 32 times while linearly varying the phase of the displacement data and the mean of those results was used as the final stiffness map. The three components of the curl were inverted

separately and then combined into a single stiffness map by computing the weighted average using the squared amplitude of each component. NNI stiffness was summarized by computing the mean within the same ROI as above (i.e., eroded by one voxel).

To compare test-retest repeatability between the DI and NNI methods, we used previously acquired data from 10 healthy volunteers who were each scanned 3 times (34). Repeatability was assessed by computing the coefficient of variation (CV) for each subject in all ROIs, and this analysis was performed using two sets of ROIs. First we used bilateral ROIs as had been previously reported including the 4 lobes of the brain and a deep gray and white matter ROI. Second, in order to investigate the effect of reduced ROI size, we used unilateral ROIs by treating the left and right hemispheres separately. To test for differences in the CV between the two inversions, the difference in CV was computed (for each subject and ROI) and then fit with a mixed-effects model that included fixed-effects for the mean difference in CV across all regions and categorical regional-effects, as well as random slopes and intercepts for each subject. A t-test was then performed on the constant fixed-effect term with a p value of less than 0.05 considered significant.

To compare the ability of DI and NNI to detect the established effect of aging on brain stiffness, we used data from 46 subjects with an equal distribution of male and female participants and age ranging from 56 to 89 (mean: 74.2 years, standard deviation: 9.15 years) (11). First, global stiffness was computed for each subject using both methods across the entire cerebrum. Least squares regression analysis was used to inspect the relationship between DI and NNI summary results using a restricted cubic spline model with 3 knots. The relationship between each stiffness estimate and age was also fit by a linear model. Finally, the effects of age and sex were modeled in the brain at the voxel-level. SPM12 was used to normalize the stiffness maps of each subject to MNI space. The T2-weighted magnitude of the MRE data were used as the reference image and this transformation was then applied to the stiffness map using nearest neighbor interpolation to preserve edges. The warped stiffness maps maintained the resolution of the original images (3-mm isotropic voxels). To reduce the amount of missing data for subsequent model fitting, the stiffness maps were then smoothed with a 3×3×3 median filter (but excluding voxels outside the brain from the calculation), and the original stiffness estimates were re-imposed for voxels within the brain (effectively dilating stiffness maps without smoothing values inside the brain). Finally, we applied a smoothed brain mask computed by performing morphological closing on the original mask. The data at each voxel (where no more than half of the data were missing) were then fit by a general linear model with predictors including age, sex, and a constant. From this model, we computed the estimated mean stiffness map across all subjects at age 75, as well as maps of the t-statistic for the effects of age and sex. These statistical maps were thresholded at a cluster-level family-wise error (FWE) corrected $P < 0.05$. This cluster-level correction (i.e., minimum cluster size) was determined using an approximate permutation test with 1,000 random permutations. More specifically, for each permutation, the subjects were randomly re-labeled before model fitting. The resulting t-statistic maps were then thresholded at $P < 0.01$, and the maximum cluster size in either statistical image (one each for the age and sex effects) was measured using an 18-connected neighborhood. The 95th percentile of the distribution of maximum cluster size then represents the minimum cluster size necessary such that the probability of finding at least

one cluster in either of the two parametric maps of interest due to chance is the specified error rate of 0.05.

Results

Simulation experiments

The results of the simulation experiments are summarized in Figure 2. In this simplest case (noise-free, single-source), both DI and NNI are shown to provide nearly perfect stiffness estimation in the test examples. This experiment represents an important proof of principle for NNI, establishing that it can accurately compute stiffness from data of varied wavelength and with arbitrary directions of propagation. When multiple sources were added, again both the DI and the NNI produced stiffness estimates that were tightly correlated with the known stiffness of the simulation. Of note, the variance of the NNI estimates appears more uniform across the range of stiffness values.

The performance in simulation of DI and NNI diverges more drastically when noise is added to the simulated data. In both the cases of single- and multiple-source simulations, the NNI estimates explain approximately $2\times$ the variance explained by DI estimates. However, non-linearity is observed in the NNI estimates near the top of the training range, driven in particular by low SNR examples.

Liver experiment

Example stiffness maps using both MMDI and NNI are shown in Figure 3, including participants with soft (left), moderate (center), and firm (right) livers as assessed by MMDI. An analogous figure with unsmoothed stiffness maps is shown in Supporting Figure S1. The summary data from all 86 subjects are shown in Figure 4. The regression analysis indicates a tight relationship between the two estimates, though the absolute values of the estimates are not precisely matched. There is significant evidence of a nonlinear relationship, as the NNI values begin to plateau for high stiffness cases (Fig. 4A). Both MMDI (Fig. 4B) and NNI (Fig. 4C) detect significant increases in stiffness with increasing fibrosis stage.

Brain experiment

The results from the test-retest repeatability study are summarized in Figure 5. Considering the bilateral ROIs, the mean CV using DI was 1.44% while the mean CV was reduced to 0.70% using NNI. Considering the unilateral ROIs, the mean CVs using both DI and NNI were 2.02% and 0.92%, respectively. In both cases, a significant difference in CV between the inversion methods was observed.

Example NNI stiffness maps in the brain for 3 subjects are shown in Figure 6. Whole-brain summary measurements are shown in Figure 7. As in the liver, we observe a strong correlation between the NNI and DI estimates of stiffness, with the NNI estimates on average 0.43 kPa greater than the DI estimates (Fig. 7A). Also, both inversions detect significant decreases in cerebral stiffness with increasing age (Fig. 7B and 7C), indicating that both are able to detect biologically relevant signals.

The results of voxel-wise modeling are shown in Figure 8, including the estimated stiffness map at age 75 in the top row. We can observe higher stiffness values in the deep gray nuclei and large white matter tracts relative to the surrounding tissue. Some artifacts are also observed including apparent overestimates at the cortical surface (likely due to the filtering method employed to reduce missing data), and apparent underestimates at the midline in superior slices (likely due to phase discontinuities created by the falx). The age-effect is summarized in the bottom row of Figure 8 (FWE corrected $P < 0.05$). Five clusters (3,343 voxels) were detected at this significance level, largely co-localizing with white matter tracts. Images of the sex-effect are not shown as no significant clusters were detected. Analogous images of the age-effect using DI and NNI, each without and with additional smoothing, are shown in Supporting Figure S2. Using unsmoothed DI stiffness maps, we observe outliers in stiffness and the detectable significant age effect is substantially reduced as compared to the NNI result (405 voxels in 13 clusters). If the DI stiffness maps are first smoothed with a $3 \times 3 \times 3$ median filter before voxel-wise modeling, the number of outliers is reduced and the age effect becomes apparent (6,203 voxels in 3 clusters). For completeness, the results with NNI stiffness maps that were similarly smoothed are also included, again detecting a substantial effect of age (8,247 voxels in 1 cluster).

Discussion

In this study, we investigated the potential to use artificial neural networks to perform inversion of MRE data. In simulation, NNI performed at least as well or better than DI for accurately predicting the known stiffness. In particular, NNI outperformed DI in the presence of noise. In liver and brain data, we showed that the mean stiffness computed by NNI was significantly correlated with stiffness estimates obtained by a conventional inversion algorithm. However, while the NNI and DI estimates were correlated *in vivo*, there was a bias between the two measurements. In the brain in particular, estimation of the Laplacian in small windows in the presence of noise typically leads to an underestimate of the actual stiffness. So the larger estimates provided by the neural network may be preliminary evidence that these stiffness measurements are less biased. Future studies in samples of known stiffness, such as in phantoms, will be necessary to evaluate the relative accuracy of these methods.

Results in the simulation experiments suggested that NNI provides greater accuracy in the presence of noise compared to an algebraic DI approach. In turn, we showed that test-retest repeatability for measuring regional brain stiffness was improved by NNI. Even when the ROIs were reduced to approximately half the original size, measurement error remained below that of DI with bilateral ROIs, demonstrating how this resistance to noise can be traded for improved measurement resolution. Taking advantage of the stability of NNI in the presence of noise, we modeled the effects of age and sex on brain stiffness at the voxel-level. We observed larger stiffness values in white matter tracts, consistent with previous literature that stiffness is associated with high fractional anisotropy as assessed by diffusion imaging (35). Stiffness of the deep gray nuclei was also greater than surrounding tissue, consistent with previous findings from our group (34), as well as Johnson et al.(36). Also consistent with previous ROI-based analyses, we observed an inverse relationship between age and stiffness throughout most of the brain, and small estimated effects due to sex (11,12). The

estimated voxel-wise age effect was substantially diminished when using the DI stiffness maps, though the effect became more apparent by applying additional smoothing to the stiffness maps. Such additional smoothing will decrease resolution, and further increase autocorrelation in these statistical maps, in turn requiring a more stringent criterion for hypothesis testing. It is possible that voxel-wise modeling of the DI stiffness maps by minimizing the absolute error (as opposed to squared error) could improve the performance; however these models would need to be solved iteratively for every voxel in the volume and therefore may be impractical. Taken together, these voxel-wise analyses suggest that NNI is more sensitive to biological signals than DI at a given resolution, but this finding will require further corroboration.

Preliminary evidence suggests that NNIs possesses certain advantages over traditional inversions, such that with further development and validation, NNIs might be a better choice in certain MRE applications. DI algorithms directly solve the equations of motion in a viscoelastic material (15–17), typically following some assumptions such as local homogeneity. DI is theoretically capable of handling interfering wave patterns, and the spatial footprint is well understood and easily controlled. However, a major drawback of this algorithm is numerical instability at the voxel-level, which must be stabilized by additional filtering of the displacement data or summarizing the stiffness maps over large ROIs. Both of these approaches substantially limit resolution. Local frequency estimation (LFE) can be used to compute stiffness by applying a filter bank to the displacement data and computing an estimate of spatial frequency from the proportion of energy in each of those filter bands (18–20). LFE typically provides stable, smooth stiffness maps, but with low resolution and a bias toward underestimation in the presence of noise. Non-linear inversion (NLI) can theoretically get around some of resolution and noise sensitivity limits that face DI and LFE because it evaluates the forward model and does not assume local homogeneity (21–23). However, the iterative approach currently used in NLI, and the requirement for multiple finite element modeling simulations within the algorithm, can lead to long reconstruction times on the order of hours. These long times may pose no problem in a research setting but are infeasible in some potential clinical applications. The NNIs for *in vivo* data presented here took approximately 8 hours to train (including generation of simulated data), and only a few minutes to reconstruct each stiffness map. This reconstruction time is expected to be even faster with a compiled GPU implementation.

Some limitations of this study will serve as the subject of future investigation. First, the simplicity of the simulated data used for training imposes some strict assumptions on the inversion. Specifically, the current model assumes tissue homogeneity in each patch, and also cannot provide information about attenuation and associated viscoelastic tissue parameters. The use of more advanced simulations and more realistic MRI-like noise may further improve accuracy *in vivo*, and allow for NNI estimation of additional mechanical parameters (e.g., storage and loss moduli or damping ratio). Second, this work uses a feedforward network architecture with fully-connected layers, which may not take full advantage of spatial information in the displacement images. Implementation of a convolutional network may better leverage this spatial information and allow for more efficient training of deeper nets. Finally, further work is necessary to understand the effects of edges on the NNIs, and how to best impute the missing data. The approach used in this

study appears to result in overestimated stiffness at edges. Nonetheless older subjects, who have more atrophy and presumably more edge-related bias, still exhibit decreasing stiffness even near the cortical surface. We expect that more accurate approaches to data imputation will provide improved sensitivity to physiological processes.

Conclusions

In conclusion, this work presents initial attempts to train artificial neural networks to estimate stiffness from images of tissue displacement obtained by MRE. In simulation experiments, the NNI is shown to perform as well or better than a standard direct inversion approach. In both liver and brain applications, NNI is shown to provide stiffness estimates that are tightly correlated with measurements made by standard approaches in the field. Particularly evident in the brain, the NNI stiffness maps were quite stable, which reduced test-retest error and enabled voxel-wise statistical modeling to produce maps of each parameter of interest. Taken together, this study demonstrates the need for further investigation into the use of artificial neural networks to estimate stiffness from MRE data.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by the National Institutes of Health grant R37-EB001981.

References

1. Muthupillai R, Lomas DJ, Rossman PJ, Greenleaf JF, Manduca A, Ehman RL. Magnetic resonance elastography by direct visualization of propagating acoustic shear waves. *Science*. 1995; 269(5232): 1854–1857. [PubMed: 7569924]
2. Asbach P, Klatt D, Hamhaber U, Braun J, Somasundaram R, Hamm B, Sack I. Assessment of liver viscoelasticity using multifrequency MR elastography. *Magnetic resonance in medicine*. 2008; 60(2):373–379. [PubMed: 18666132]
3. Huwart L, Sempoux C, Salameh N, Jamart J, Annet L, Sinkus R, Peeters F, ter Beek LC, Horsmans Y, Van Beers BE. Liver fibrosis: Noninvasive assessment with MR elastography versus aspartate aminotransferase-to-platelet ratio index. *Radiology*. 2007; 245(2):458–466. [PubMed: 17940304]
4. Rustogi R, Horowitz J, Harmath C, Wang Y, Chalian H, Ganger DR, Chen ZME, Bolster BD, Shah S, Miller FH. Accuracy of MR elastography and anatomic MR imaging features in the diagnosis of severe hepatic fibrosis and cirrhosis. *J Magn Reson Imaging*. 2012; 35(6):1356–1364. [PubMed: 22246952]
5. Yin M, Glaser KJ, Talwalkar JA, Chen J, Manduca A, Ehman RL. Hepatic MR Elastography: Clinical Performance in a Series of 1377 Consecutive Examinations. *Radiology*. 2016; 278(1):114–124. [PubMed: 26162026]
6. ElSheikh M, Arani A, Perry A, Boeve BF, Meyer FB, Savica R, Ehman RL, Huston J 3rd. MR Elastography Demonstrates Unique Regional Brain Stiffness Patterns in Dementias. *AJR American journal of roentgenology*. 2017:1–6.
7. Murphy MC, Jones DT, Jack CR Jr, Glaser KJ, Senjem ML, Manduca A, Felmlee JP, Carter RE, Ehman RL, Huston J 3rd. Regional brain stiffness changes across the Alzheimer's disease spectrum. *NeuroImage Clinical*. 2016; 10:283–290. [PubMed: 26900568]

8. Streitberger KJ, Sack I, Krefting D, Pfuller C, Braun J, Paul F, Wuerfel J. Brain viscoelasticity alteration in chronic-progressive multiple sclerosis. *PloS one*. 2012; 7(1):e29888. [PubMed: 22276134]
9. Streitberger KJ, Wiener E, Hoffmann J, Freimann FB, Klatt D, Braun J, Lin K, McLaughlin J, Sprung C, Klingebiel R, Sack I. In vivo viscoelastic properties of the brain in normal pressure hydrocephalus. *NMR in biomedicine*. 2011; 24(4):385–392. [PubMed: 20931563]
10. Wuerfel J, Paul F, Beierbach B, Hamhaber U, Klatt D, Papazoglou S, Zipp F, Martus P, Braun J, Sack I. MR-elastography reveals degradation of tissue integrity in multiple sclerosis. *NeuroImage*. 2010; 49(3):2520–2525. [PubMed: 19539039]
11. Arani A, Murphy MC, Glaser KJ, Manduca A, Lake DS, Kruse SA, Jack CR Jr, Ehman RL, Huston J 3rd. Measuring the effects of aging and sex on regional brain stiffness with MR elastography in healthy older adults. *NeuroImage*. 2015; 111:59–64. [PubMed: 25698157]
12. Sack I, Beierbach B, Wuerfel J, Klatt D, Hamhaber U, Papazoglou S, Martus P, Braun J. The impact of aging and gender on brain viscoelasticity. *NeuroImage*. 2009; 46(3):652–657. [PubMed: 19281851]
13. Schwarb H, Johnson CL, Daugherty AM, Hillman CH, Kramer AF, Cohen NJ, Barbey AK. Aerobic fitness, hippocampal viscoelasticity, and relational memory performance. *NeuroImage*. 2017; 153:179–188. [PubMed: 28366763]
14. Schwarb H, Johnson CL, McGarry MD, Cohen NJ. Medial temporal lobe viscoelasticity and relational memory performance. *NeuroImage*. 2016; 132:534–541. [PubMed: 26931816]
15. Hirsch S, Guo J, Reiter R, Papazoglou S, Kroenke T, Braun J, Sack I. MR Elastography of the Liver and the Spleen Using a Piezoelectric Driver, Single-Shot Wave-Field Acquisition, and Multifrequency Dual Parameter Reconstruction. *Magnetic resonance in medicine*. 2014; 71(1): 267–277. [PubMed: 23413115]
16. Oliphant TE, Manduca A, Ehman RL, Greenleaf JF. Complex-valued stiffness reconstruction for magnetic resonance elastography by algebraic inversion of the differential equation. *Magnetic resonance in medicine*. 2001; 45(2):299–310. [PubMed: 11180438]
17. Papazoglou S, Hamhaber U, Braun J, Sack I. Algebraic Helmholtz inversion in planar magnetic resonance elastography. *Phys Med Biol*. 2008; 53(12):3147–3158. [PubMed: 18495979]
18. Braun J, Buntkowsky G, Bernarding J, Tolxdorff T, Sack I. Simulation and analysis of magnetic resonance elastography wave images using coupled harmonic oscillators and Gaussian local frequency estimation. *Magn Reson Imaging*. 2001; 19(5):703–713. [PubMed: 11672629]
19. Clayton EH, Okamoto RJ, Bayly PV. Mechanical Properties of Viscoelastic Media by Local Frequency Estimation of Divergence-Free Wave Fields. *J Biomech Eng-T Asme*. 2013; 135(2)
20. Manduca A, Oliphant TE, Dresner MA, Mahowald JL, Kruse SA, Amromin E, Felmlee JP, Greenleaf JF, Ehman RL. Magnetic resonance elastography: non-invasive mapping of tissue elasticity. *Med Image Anal*. 2001; 5(4):237–254. [PubMed: 11731304]
21. McGarry MDJ, Van Houten EEW, Johnson CL, Georgiadis JG, Sutton BP, Weaver JB, Paulsen KD. Multiresolution MR elastography using nonlinear inversion. *Med Phys*. 2012; 39(10):6388–6396. [PubMed: 23039674]
22. Van Houten EE, Miga MI, Weaver JB, Kennedy FE, Paulsen KD. Three-dimensional subzone-based reconstruction algorithm for MR elastography. *Magnetic resonance in medicine*. 2001; 45(5):827–837. [PubMed: 11323809]
23. Van Houten EE, Paulsen KD, Miga MI, Kennedy FE, Weaver JB. An overlapping subzone technique for MR-based elastic property reconstruction. *Magnetic resonance in medicine*. 1999; 42(4):779–786. [PubMed: 10502768]
24. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015; 521(7553):436–444. [PubMed: 26017442]
25. Duda, RO., Hart, PE., Stork, DG. Pattern classification. New York: Wiley; 2001. p. 654xx
26. Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *J Mach Learn Res*. 2014; 15:1929–1958.
27. Rumelhart DE, Hinton GE, Williams RJ. Learning Representations by Back-Propagating Errors. *Nature*. 1986; 323(6088):533–536.

28. Moller MF. A Scaled Conjugate-Gradient Algorithm for Fast Supervised Learning. *Neural Networks*. 1993; 6(4):525–533.
29. Chen J, Yin M, Talwalkar JA, Oudry J, Glaser KJ, Smyrk TC, Miette V, Sandrin L, Ehman RL. Diagnostic Performance of MR Elastography and Vibration-controlled Transient Elastography in the Detection of Hepatic Fibrosis in Patients with Severe to Morbid Obesity. *Radiology*. 2017; 283(2):418–428. [PubMed: 27861111]
30. Silva AM, Grimm RC, Glaser KJ, Fu Y, Wu T, Ehman RL, Silva AC. Magnetic resonance elastography: evaluation of new inversion algorithm and quantitative analysis method. *Abdominal imaging*. 2015; 40(4):810–817. [PubMed: 25742725]
31. Manduca A, Lake DS, Kruse SA, Ehman RL. Spatio-temporal directional filtering for improved inversion of MR elastography images. *Med Image Anal*. 2003; 7(4):465–473. [PubMed: 14561551]
32. Okamoto RJ, Clayton EH, Bayly PV. Viscoelastic properties of soft gels: comparison of magnetic resonance elastography and dynamic shear testing in the shear wave regime. *Phys Med Biol*. 2011; 56(19):6379–6400. [PubMed: 21908903]
33. Coifman, RR., Donoho, DL. Translation-Invariant De-Noising. In: Antoniadis, A., Oppenheim, G., editors. *Wavelets and Statistics*. New York, NY: Springer New York; 1995. p. 125-150.
34. Murphy MC, Huston J 3rd, Jack CR Jr, Glaser KJ, Senjem ML, Chen J, Manduca A, Felmlee JP, Ehman RL. Measuring the characteristic topography of brain stiffness with magnetic resonance elastography. *PloS one*. 2013; 8(12):e81668. [PubMed: 24312570]
35. Johnson CL, McGarry MD, Gharibans AA, Weaver JB, Paulsen KD, Wang H, Olivero WC, Sutton BP, Georgiadis JG. Local mechanical properties of white matter structures in the human brain. *NeuroImage*. 2013; 79:145–152. [PubMed: 23644001]
36. Johnson CL, Schwarb H, M DJM, Anderson AT, Huesmann GR, Sutton BP, Cohen NJ. Viscoelasticity of subcortical gray matter structures. *Human brain mapping*. 2016; 37(12):4221–4233. [PubMed: 27401228]

Figure 1.

A. Model neuron used in an artificial neural network. The neuron computes the sum of the product of weights, w , and inputs, x , as well as a bias term, and then applies a nonlinear transfer function, h . B. Diagram of an example feedforward neural network with 3 inputs, 2 hidden layers with 5 units in each, and a single output.

Figure 3.

Example images in the liver. Each column shows a single slice of data from three subjects with progressively increasing stiffness. The first row is a single phase offset of the wave images. Stiffness maps calculated by MMDI are shown in the second row and maps calculated by NNI are shown in the third row. The region of interest for computing the mean stiffness is outlined in white in each case.

Figure 4.

Summary results in the liver. A. Comparison of MMDI and NNI results. Data were fit by a restricted cubic spline model with 3 knots. The blue shaded region indicates the 95% confidence interval calculated from 1,000 bootstrapped samples. B. Summary of MMDI stiffness as a function of fibrosis stage. C. Summary of NNI stiffness as a function of fibrosis stage. For each group, boxes represent the interquartile range (IQR) with the red lines indicating the median stiffness. Whiskers depict the full range of the data, excluding outliers ($>1.5 \times \text{IQR}$ from the median), which are indicated by red plus signs.

Figure 5.

Summary of test-retest repeatability using both DI and NNI in the brain. Coefficient of variation (CV) was computed across 3 trials for each subject in each ROI. The graph displays the mean and standard deviation of the measured CV over the 10 volunteers. Two sets of ROIs were considered. Bilateral ROIs included the lobes of the brain and a ROI including the deep gray matter and white matter. Unilateral ROIs cut the ROI size approximately in half by treating the left and right hemispheres separately. Differences in the CV were tested using a mixed-effects model with a t-test on the term describing the difference in mean CV across all regions and subjects ($p < 0.05$ considered significant).

Figure 6.
Example images in the brain. Each row shows the stiffness map for one subject at 5 slice locations arranged from inferior to superior positions. The age of the subjects increases from top to bottom.

Figure 7.

Summary results for global brain stiffness. A. Comparison of DI and NNI results. B. Relationship between age and DI stiffness. C. Relationship between age and NNI stiffness. In each panel, data were fit by a regression model (Panel A: restricted cubic spline model with 3 knots; Panels B and C: linear model). The blue shaded region indicates the 95% confidence interval calculated from 1,000 bootstrapped samples.

Figure 8.

Summary of voxel-wise modeling. Ten slice locations are shown, arranged from inferior to superior positions. The estimated stiffness map at age 75 obtained by modeling data from all 46 subjects using NNI stiffness maps is shown in the top row. In the bottom row is a map of the effect of age. Results are thresholded at a family-wise error corrected $P < 0.05$ as determined by a permutation test ($P < 0.01$ primary threshold and cluster size > 82 voxels). The resulting thresholded T-map is overlaid on a grayscale mean stiffness map.