Published in final edited form as:

Biomech Model Mechanobiol. 2017 June ; 16(3): 805–822. doi:10.1007/s10237-016-0854-6.

An Information-Based Machine Learning Approach to Elasticity Imaging

Cameron Hoerig,

Department of Civil and Environmental Engineering, University of Illinois at Urbana-Champaign, Urbana, IL

Jamshid Ghaboussi, and

Department of Civil and Environmental Engineering, University of Illinois at Urbana-Champaign, Urbana, IL

Michael. F. Insana

Department of Bioengineering and Beckman Institute for Advanced Science and Technology, University of Illinois at Urbana-Champaign, Urbana, IL

Abstract

An information-based technique is described for applications in mechanical-property imaging of soft biological media under quasi-static loads. We adapted the Autoprogressive method that was originally developed for civil engineering applications for this purpose. The Autoprogressive method is a computational technique that combines knowledge of object shape and a sparse distribution of force and displacement measurements with finite-element analyses and artificial neural networks to estimate a complete set of stress and strain vectors. Elasticity imaging parameters are then computed from estimated stresses and strains. We introduce the technique using ultrasonic pulse-echo measurements in simple gelatin imaging phantoms having linear-elastic properties so that conventional finite-element modeling can be used to validate results. The Autoprogressive algorithm does not require any assumptions about the material properties and can, in principle, be used to image media with arbitrary properties. We show that by selecting a few well-chosen force-displacement measurements that are appropriately applied during training and establish convergence, we can estimate all nontrivial stress and strain vectors throughout an object and accurately estimate an elastic modulus at high spatial resolution. This new method of modeling the mechanical properties of tissue-like materials introduces a unique method of solving the inverse problem and is the first technique for imaging stress without assuming the underlying constitutive model.

Keywords

Finite-element Analysis; Mechanical Properties; Neural Networks; Ultrasonic Speckle Tracking; Constitutive Modeling

1 Introduction

Medical elasticity imaging encompasses a broad range of techniques for imaging mechanical properties of biological tissues. Standard approaches to patient imaging begin by tracking displacements resulting from a weak mechanical stimulus applied to the tissue being imaged. Combining time-varying displacement information with measurements of or assumptions about the associated forces into a collection of linear equations, the elements of a constitutive matrix are estimated.

Elasticity imaging reveals how tissue stiffness and/or viscosity vary with position, time, and applied force. These basic mechanical properties can indicate regions of inflammation, edema, hypertrophy, and fibrosis that accompany the presence of disease processes (Barr 2014). Conventional approaches show great promise for cancer imaging in particular. However, simplifying assumptions, which enter the analysis as tissue models, are often unjustified given the structural complexity of most tissues. These assumptions are necessary to formulate a properly constrained inverse problem for mechanical parameter estimation. Violating the assumptions of a tissue model can distort the description of the mechanical environment needed for diagnostic decision making. Unjustified assumptions also limit the vast possibilities for learning more about the role of mechanobiology in revealing disease processes. Furthermore, simple models in common use generally discard some of the information gathered during an imaging exam.

This report describes a different approach to elasticity imaging that requires no assumptions about tissue mechanical properties, including linearity and isotropy. Our initial focus is on quasi-static methods that generate full stress and strain maps from just a few force-displacement measurements and knowledge of object shape. Estimation of constituent-matrix elements for elasticity imaging occurs retrospectively, once the stress and strain fields have been accurately estimated.

The technique accepts a time-series of surface force-displacement measurements made while pressing a rigid ultrasound transducer into the medium surface. Simultaneously, the transducer records radiofrequency echo data used to estimate a time series of displacements at select points in the medium. These force-displacement data are input into the Autoprogressive (AutoP) method (Ghaboussi et al 1998) that employs two finite-element analyses (FEAs) to exploit (a) equilibrium conditions relating forces and stresses and (b) compatibility requirements relating displacements and strains. We will show how these FEA processes occur simultaneously and iteratively by relating them to each other through one or more artificial neural networks (ANNs) that learn and record material properties. The ANNs take the place of the constitutive matrix, relieving us of the need for simplifying assumptions. The material properties learned by the ANNs during training with AutoP are stored in the distributed connection weights. Later, and without further measurements, the informational model can be probed to find imaging parameters in the range of mechanical stimuli that were used to train the model.

An important part of quasi-static deformation analysis is that a mechanical stimulus applied to the medium may be sensed at every contiguous location in that medium. Consequently, a

few well-positioned measurements made during quasi-static deformation can result in a complete mechanical description, even for large deformations in nonlinear media. Combining FEA and ANNs in AutoP allows the material properties to be extracted from a sparse sampling of force-displacement measurements.

For the first time, complete and accurate images of all stress and strain fields throughout an object resulting from an applied force can be described from experimental measurements. The advancement to the art is in combining a few strategically-placed measurements with mechanical principles and information-based machine learning tools to solve the inverse problem without making any initial assumptions of the underlying constitutive model. Furthermore, the trained ANNs are not used to classify or interpret images; instead, they are used to construct images of material properties. Our informational models usher in a new approach to implementing ANNs for medical imaging applications.

Despite some similarities with previous reconstruction methods in elasticity imaging that use FEA to iteratively solve the inverse problem (Kallel and Bertrand 1996; Tyagi et al 2014), we do not apply mathematical inverse methods to directly estimate a constitutive matrix. Existing iterative schemes still require the assumptions of a constitutive model with a compact feature set. The ANNs in AutoP replace the constitutive matrix, eliminating the need for assumptions about material properties and boundary values.

In this introductory report, we consider relatively simple physical phantoms with known properties to demonstrate the method. Heterogeneous phantoms composed of linear-elastic gelatin gels with known shapes and stiffnesses were built to develop and validate these methods. An ultrasonic linear array is used to compress phantoms under plane-stress conditions. We expand upon our previous reports (Hoerig et al 2015, 2016) and show that all stress and strain fields for 2-D and 3-D measurement geometries, along with the spatially varying Young's modulus, can be accurately reconstructed even for locations outside the region sampled ultrasonically.

2 Methods

A constitutive model relates output stress to input strain (or vice-versa) through mathematical equations based on known physical principles and experimentally observed behavior. For systems too complex to have a complete reductionist description, informational models can describe input-output relationships via machine learning techniques and training involving repeated exposure to high-quality experimental data. In the following section, we describe AutoP as applied to data from tissue-like gelatin phantoms.

2.1 Force-displacement Measurements

Imaging phantoms were constructed using a simple gelatin formulation of water, gelatin powder, and corn starch that acts as a scattering agent (Orescanin and Insana 2010). Each phantom was a $50 \times 50 \times 50 \text{ mm}^3$ cube of a soft gel ($\approx 9 \text{ kPa}$) with one or three embedded stiff cylindrical inclusions ($\approx 22\text{--}26 \text{ kPa}$), 10 mm in diameter and 50 mm long. Dimensions of the phantoms and placement of the inclusions were tightly controlled via machined molds. By pouring the gelatin mixture for the inclusions after the background material solidified, we

were able to create a tightly coupled bond between the background and inclusions. Macro-indentation methods (Altahhan et al 2016) were applied to independently estimate the Young's modulus of each gelatin material. The long axis of each inclusion was oriented along the x_3 axis (out of plane of Fig. 1). Each phantom component is a linear-elastic, incompressible material with a gelatin powder concentration that determines the gel stiffness.

The basic measurement technique involves pressing an ultrasound linear-array transducer into the phantom surface in a series of discrete steps (Fig. 1). The US probe face was positioned flush with the top surface of the phantom in the (x_1, x_3) plane that are, respectively, the lateral and elevational axes of the transducer. A plastic housing used for positioning encased the US probe, creating a phantom contact area of $43 \times 50 \text{ mm}^2$. A 6-axis force-torque transducer (ATI Industrial Automation, Apex, NC) coupled the ultrasound probe housing to a motion controller capable of positioning the probe with sub-millimeter accuracy (not shown). The bottom surface of the phantom was fixed to a rigid base so it did not slip as the probe compressed the top phantom surface. Coupling gel was applied to the top surface and provided free-slip contact between the probe and phantom. The remaining four surfaces of the phantom were unconstrained during loading.

In this study, only a downward compressive force was slowly applied along the axis of the ultrasound beam (x_2 axis). A radiofrequency (RF) echo-signal frame was acquired after each load increment using a Siemens Sonoline Antares ultrasound system (Siemens Healthcare USA, Mountainview, CA) with a VF10-5 linear-array probe at 8 MHz center frequency. At the same time, the axial surface force generated by the US probe was measured by the force-torque transducer and the probe position recorded via the motion controller. Application of a speckle-tracking algorithm (Rivaz et al 2011) provided an estimate of internal displacements within the RF frames (see Appendix 1 for an evaluation of the accuracy of the speckle-tracking algorithm). One series of load increments provided a set of surface forces and probe displacements that were measured directly and a corresponding set of displacements within the phantoms which required speckle tracking to estimate. These data were then used as inputs to AutoP to build informational models of the gelatin phantoms. Information regarding the geometry of the phantoms was also necessary to develop finite-element meshes used in AutoP. External boundaries were known, but the shape and location of internal structures must be assumed or estimated. In Sections 2.4 and 2.5, we investigate cases where the inclusion boundaries were modeled exactly or were estimated via manual segmentation.

2.2 Autoprogressive Method

2.2.1 Overview—At the core of our approach to create empirical models of mechanical behavior is the Autoprogressive Method developed by Ghaboussi et al (1998). The power of AutoP lies in its ability to generate stress-strain information from force-displacement measurements. AutoP makes use of standard FEA methods (ABAQUS 6.13 finite-element software) to iteratively generate stress and strain data as artificial neural networks learn material properties.

A finite-element (FE) mesh is constructed that models the data acquisition process. For our measurements described in later sections, the FE mesh comprised a rigid, rectangular model of the US probe in contact with a square (or cube for 3-D) mesh of a phantom. Then, measured forces and displacements are split and applied to the FE model to solve two different FEAs: FEA_{σ} and FEA_{ϵ} . Solving FEA_{σ} allows us to estimate the stresses throughout the model in response to an applied force. Similarly, solving FEA_{ϵ} provides an estimate of the strains resulting from an applied displacement. Because forces and displacements are both measured for each experiment, the strains provided by FEA_{ϵ} should be consistent with the stresses found in FEA_{σ} when the ANNs that relate the two FEAs correctly represent material properties. Iteratively solving FEA_{σ} and FEA_{ϵ} while adjusting ANN weights produces increasingly accurate estimates of the stress and strain existing in the medium at the time forces and displacements were measured. Through the course of AutoP, the ANNs learn the material properties based on the stress-strain behavior estimated by the FEAs.

Understanding how AutoP works requires exploring how the ANNs affect the the FEAs (and vice-versa) and the underlying problem being solved. AutoP does not merely use FEA and ANNs as tools; rather, the two are tightly coupled so that solutions to FEA_{σ} and FEA_{ϵ} are dependent on the current state of the ANN material model(s). Both FEAs are solving the problem

$$\mathbf{K}\mathbf{U}=\mathbf{P} \quad (1)$$

where $\mathbf{K}$ is an $N \times N$ matrix constructed based on the material properties described by the ANNs, $\mathbf{U}$ is a $N \times 1$ vector of displacements for each node in the mesh, and $\mathbf{P}$ is a $N \times 1$ vector of forces. The ANNs are involved during the calculation of 1) the internal forces of an element ($\mathbf{P}$) in the FE mesh and 2) the Jacobian matrix (i.e., consistent tangent). The integration required to compute the internal forces of an element contains a stress term that varies based on strain, a term that is dependent on displacement ($\mathbf{U}$). Providing the strain as input to the ANN(s) produces the required stress term.

Computing the Jacobian matrix - which is a necessary part of the Newton-Raphson iterations utilized in non-linear FEA solvers - is not straightforward. A thorough explanation of how the Jacobian matrix is computed can be found in (Hashash et al 2004), but it should be noted that the Jacobian has a large impact on the FEA solution. Progressing through AutoP iteratively updates the connection weights of the ANN(s), which in turn changes the computed internal forces and Jacobian matrix, producing different stress and strain distributions and solutions of $\mathbf{P}$ and $\mathbf{U}$ in FEA_{σ} and FEA_{ϵ} . It is this feedback loop - the change in the ANN affecting the computed stresses and strains that in turn updates the ANN - that allows the ANN(s) to gradually learn the mechanical properties that generate the stresses and strains consistent with the force-displacement measurements. If the ANNs perfectly characterize the mechanical properties of the modeled object (i.e., $\mathbf{K}$ is correct), the solutions for $\mathbf{U}$ and $\mathbf{P}$ from both FEA_{σ} and FEA_{ϵ} will match.

ANNs are selected as the machine learning component of AutoP because of their inherent nonlinearity, robustness, and ability to learn complex behaviors (Ghaboussi 2010; Ghaboussi and Sidarta 1998). Their highly parallel structure accumulates information as it becomes available experimentally. A general concern when using ANNs is determining the size of the network to be used. ANNs with too many nodes in each layer have a capacity exceeding that necessary to learn mechanical behavior and can generate overfitting errors. On the other hand, if there are too few nodes per layer, the ANN does not have enough capacity to store the necessary information. We implemented nested adaptive neural networks (Ghaboussi and Sidarta 1998) to minimize these issues.

2.2.2 Creating Finite-element Meshes—Training with AutoP begins with the development of a FE model of the data acquisition process. In the FE model, probe-phantom contact and external boundary conditions are set to mimic the experimental setup; for compressions of the gelatin phantoms, the bottom edge of the phantom is fixed, the contact surface between the ultrasound probe and phantom is assumed frictionless, and the remaining sides are free. In this initial technique, one ANN is used for each region that could have distinct material properties. Here, only two ANNs were used: ANN_{soft} which describes the material properties of the soft background gelatin material, and ANN_{stiff} which characterizes the stiff inclusion(s). Only one ANN is needed for the inclusions in the three-inclusion model because each is composed of the same gelatin mixture. If a different material was used for each inclusion, three different ANNs would be required, bringing the total to four. Shape estimation is required for meshing and to determine the number of ANNs applied.

2.2.3 Elastic Pre-training—Once the FE mesh is created and the ANNs are assigned to the mesh elements, each ANN material model is pre-trained with stress-strain data generated using linear-elastic equations. This phase is for ANN initialization only to avoid non-physical behavior that emerges from random connection weights. Note that the model selected for pre-training is not important. During the course of training with AutoP, the linear-elastic behavior learned during pre-training is “overwritten” as the ANNs learn the mechanical behavior of a specific object. We will show that the ANNs can be pre-trained using data generated with an incorrect assumption of the Young’s modulus (i.e., different than the value estimated from macro-indentation measurements) but will learn the stress-strain behavior governing the measured force-displacement data, allowing us to accurately estimate the Young’s modulus after training with AutoP.

2.2.4 AutoP Training Iterations—AutoP training iterations, diagrammed in Fig. 2, begin after the pre-training phase. Force-displacement data acquired during the first load step (first compression of the gelatin phantom) are applied to the FE model to solve FEA _{σ} and FEA _{ϵ} . Measured forces are applied in FEA _{σ} and the FEA is solved using the current state of the ANNs ([NN] _{i}) as the element material models. In the solution of the FEA, equilibrium conditions -which relate forces and stresses - are satisfied, producing an estimate of the stresses throughout the model of the phantom. In the second analysis, FEA _{ϵ} , measured displacements are applied and the FEA solved once again using the current state of the ANNs as the element material models. Compatibility conditions, which relate displacements

to strains, are satisfied in FEA_e , producing an estimate of the strains throughout the phantom model. After solving both FEAs, stress vectors computed in FEA_σ are paired with their corresponding strain vectors computed in FEA_e to create a set of stress-strain training pairs used to update the connections weights of the ANNs ($[NM]_j \rightarrow [NM]_{j+1}$).

Training iterations are performed using the same force-displacement data from the first load step until a predefined maximum number of iterations is reached or the ANNs converge. Convergence of the ANN material models is checked by using the computed displacements from $FEA_\sigma(\hat{u}_j)$ and the measured displacements provided in $FEA_e(u_j)$ to calculate a vector of displacement errors using Eq. (2).

$$\Delta u_i = |u_i - \hat{u}_i| \quad (2)$$

Note that the total number of computed displacement errors matches the number of displacements provided in FEA_e . In Section 3.1.2, we discuss the effect of including additional displacements not used in FEA_e . These displacement errors are used to calculate two values:

$$c_{max} = \frac{\max(\Delta u_i)}{|\max(u_i)|} \quad (3)$$

$$c_\mu = \frac{\text{mean}(\Delta u_i)}{|\max(u_i)|} \quad (4)$$

If $c_{max} < C_{max}^n$ and $c_\mu < C_\mu^n$, where C_{max}^n and C_μ^n are predefined criteria for the n^{th} load step, convergence has been achieved for the first load step. AutoP iterations are repeated in the same manner using each set of force-displacement data for all load steps. This constitutes one pass of AutoP. Multiple passes may be performed using the same or new data.

The stress-strain response of an ANN during the course of training is visualized in Fig. 3. In this example, the true linear-elastic behavior of an imaginary 1-D material (—) is to be learned by an ANN that has been pretrained (---). Force-displacement data from the first load step (Load Step 1) are used in AutoP iterations, causing the stress-strain response of the ANN to gradually move toward the true behavior. Convergence is achieved after several iterations, implying the mechanical behavior described by the ANN matches the true behavior over the stress-strain range produced in the first load step (- - -). Outside of this range, the ANN response becomes nonlinear and diverges from the true response. Performing additional AutoP iterations with force-displacement measurements from other load steps extends the range over which the ANN and the true behavior match (- · - ·). Note that if the true mechanical behavior was nonlinear, the inherent nonlinear behavior of ANNs would allow this method to learn those responses from the data.

2.2.5 Mechanical Parameter Estimation—The output of AutoP is one or more trained ANNs that accurately describe the mechanical behavior of every material in the target object over the stress-strain range resulting from the applied loads (e.g., compressions by the US probe). A fully trained ANN essentially replaces the constitutive matrix to compute an output stress vector from an input strain vector as illustrated in Fig. 4. We can then conduct numerical experiments where we apply any force (or displacement) within the training range to estimate all components of stress and strain. This data may be applied to a constitutive model, e.g., Kelvin-Voigt, and estimate an elastic modulus or other mechanical parameter that compose the image. This is a profound step away from current elastography methods that require a model to be assumed at the beginning in order to derive a properly constrained inverse problem. With our method, we can explore the parameter space and determine which parameter best fits the mechanical behavior *after* the stress-strain behavior is learned by the ANNs. In section 3.1, we will show a method of interrogating a trained ANN material model to determine an appropriate constitutive model for parameter estimation.

2.3 Experimental Design

Testing of this new method of modeling the mechanical properties of materials was performed on gelatin phantoms using the data acquisition scheme described in Section 2.1. A total of four phantoms were created having background gelatin materials of nominally the same stiffness and one or three stiffer inclusions. Phantoms 1 and 4 contained a single inclusion centered in the (x_1, x_2) plane. Phantom 2 contained a single non-centrally positioned inclusion, and Phantom 3 embeds three inclusions. Example B-mode images of the four phantoms are provided in Fig. 5.

In the following sections, we will describe specific details regarding data acquisition and AutoP training for models under measurement geometries. Regardless of the model, AutoP was used to train feed-forward, fully connected ANNs with hyperbolic tangent activation functions. Each ANN contained two hidden layers and resilient propagation (Riedmiller and Braun 1993) was chosen as the method to update ANN connection weights. The adaptive ANN architecture helped alleviate issues of ANN size. The beginning and ending sizes of the hidden layers in the ANNs varied and will be specified in the following sections.

A total of 30 compressions (load increments) were performed on each phantom. Total probe displacement was 1.5 mm for Phantoms 1–3 (3% of the 50mm phantom height) and 2.5 mm for Phantom 4 (5% of phantom height). Axial forces measured between 12–29 N. The total phantom deformation was kept small in these preliminary studies to ensure linear-elastic response of the phantoms. In future studies, we will show this restriction can be lifted. We trained ANNs with AutoP using 10 nearly equally spaced load steps that spanned the entire range of the applied load, including the first and last load increments. This means force-displacement data from load increments 1, 4, 7, ..., 26, 30 were used in AutoP. Small load increments were used to minimize artifacts caused by decorrelation and out-of-plane motion when applying the speckle-tracking algorithm for internal displacement estimation. In FEA, measured forces were applied as concentrated loads to the top nodes of US probe in the FE meshes. Measured probe displacements were applied to the top nodes of probe (same nodes forces were applied) and estimated internal displacements were given at three points

within each inclusion in FEA_{σ} . For example, the highlighted points in Fig. 5 indicate the specific locations where interior displacements were applied.

In the elastic pre-training phase, a set of random stress-strain vectors were generated using the plane-stress approximation (for 2-D models) or the constitutive matrix for a Hookean solid (for 3-D models). Young's moduli that differed by at least 30% from the values estimated during macro-indentation were chosen to create the pre-training pairs. For example, the Young's modulus of the background gelatin material in Phantom 1 was estimated to be 9.16 kPa from indentation testing whereas a value of 4 kPa was chosen for pre-training ANN_{soft} . We chose values of Young's moduli that significantly differed from true values to show that our informational models are capable of producing accurate estimates of material properties when initialized to an incorrect state.

After gathering the stress-strain training pairs in an iteration of AutoP, we adjusted the data in two ways. First, the mean and standard deviation of each strain component was computed. We then removed any training pairs with a strain component falling outside of the mean $\pm$ one standard deviation range. Preliminary tests indicated that limiting the variation of new information to be learned by the ANNs improved their ability to learn by reducing the impact of outlier data during the weight update stage.

Second, frame invariance (i.e., objectivity) was enforced. Frame invariance means the stress-strain response learned by the ANNs is independent of the chosen axes; that is, if the coordinate system was rotated $\frac{\pi}{2}$ (observed from a different perspective), the resulting stress and strain distributions are unchanged. Objectivity is a fundamental concept in mechanics and requires no assumption of the underlying constitutive behavior. By swapping the normal stress and strain components in each training pair, the ANNs learn frame invariance. Enforcing frame invariance doubled (or tripled in the case of 3-D) the training data obtained in each iteration of AutoP. We must stress that frame invariance is built into the training data to ensure the ANNs learn the fundamental principle, not to increase the total amount of training data. Furthermore, the stress-strain vectors could be rotated by any arbitrary angle in range $0-2\pi$; the angle $\frac{\pi}{2}$ was chosen because it did not require any matrix multiplications, therefore reducing the computational load.

A forward FEA, identified as FEA_{ANN}^n , was performed at the end of each training step. This was effectively FEA_{σ} with the exception that all 10 force measurements were applied.

Results from FEA_{ANN}^n allowed us to retrospectively analyze the state of the ANNs during AutoP. Stresses and strains computed in FEA_{ANN}^{10} - performed upon the completion of AutoP - are of particular interest as they are produced by the fully trained ANNs and will be referenced in Section 3.

2.4 2-D Models with Known Internal Geometry

Informational models were first developed for 2-D models of Phantoms 1–3 under the assumption that the internal geometry was known. For this, the FE meshes modeled the size and locations of the inclusions determined by the phantom molds, as illustrated in Figs. 5a–5c. Each FE mesh aligns with the B-mode image within 1mm error, verified by comparing

the boundary of the inclusions in the B-mode and mesh. A plane-stress approximation was used in these 2-D models of the phantoms. Therefore, axial, lateral, and shear strains and stresses were used as the input and output of the ANNs, respectively. Four-node, quadrilateral, generalized plane-stress elements (CPS4 in ABAQUS 6.13) were used to create the FE meshes of the phantoms. The US probe was meshed using the same element type, although it was modeled as a rigid body.

Additional data were acquired on Phantoms 2 and 3 by rotating the phantoms about the x_3 axis and applying compressive loads in the same way. New FE meshes were created to model the rotated versions of the phantoms. These meshes were virtually the same as those in Figs. 5b and 5c, albeit rotated 270° or 90°, respectively.

AutoP training was performed over two passes at ten load steps each. Up to three training iterations were performed in load steps 1–3 and up to five iterations in the remaining steps. Convergence criteria (described in Section 2.2.1) were set as $C_{max}^n=0.5$ and $C_{\mu}^n=0.1$ for all load steps. Displacement errors (u_i) were computed only for those nodes provided displacement information in FEA_e (i.e., only for the nodes highlighted with dots in Fig. 5).

ANNs were initialized to have three nodes in each hidden layer. Elastic pre-training was performed by generating 100 random stress-strain vectors for an incompressible material (Poisson's ratio $\nu=0.5$) as described in Section 2.2.3. As AutoP progressed, one node was added to each hidden layer during the second training iteration in load steps 4 and 5 in the first pass, resulting in ANNs that had five nodes per hidden layer at the completion of training.

Two important aspects of our new method are that informational models are 1) better able to learn mechanical behavior when provided with diverse training data and 2) able to update their knowledge as new information becomes available. In AutoP, one way to produce a richer set of training data is to apply a variety of spatial load distributions to the same phantom. For Phantom 1, we acquired data with the US probe centrally positioned in the (x_1 , x_2) plane (data set 1, previously described) and with the probe located 4 mm off-center along the x_1 axis (data set 2). Data set 2 was acquired in the same manner as data set 1. From these two data sets on Phantom 1, we explored four training scenarios described as cases in Table 1. Cases (a) and (b) compared with cases (c) and (d) explores 1) how the diversity of training data affects modulus estimation accuracy as compared to macro-indentation testing and 2) the ability of ANN material models to incorporate new information and its affect on modulus estimates.

2.5 2-D Models with Coarsely Estimated Internal Geometry

Imaging in a clinical setting precludes prior knowledge of the exact internal geometry. Instead, images obtained during data acquisition must be segmented to determine boundaries of each unique material in the object. We tested the ability of our informational models to learn the mechanical properties of the gelatin materials when geometry errors are present in the FE models.

B-mode images of Phantoms 1 and 3 were manually segmented to create the FE meshes used for training. Two cases of segmentation were performed and are described in Table 1. In case (e), the boundaries of the inclusion were closely followed, producing a FE mesh with internal boundaries very similar to the meshes in Fig. 5. Conversely, case (f) is a segmentation producing large errors in estimating the edge of the inclusion. We note that while the shape of the inclusion was a poor representation of the true geometry in case (f), the general location of the inclusion was accurate.

Training cases (e) and (f) were performed for Phantom 1, while only case (e) was performed for Phantom 3. This latter case poses an interesting problem. The B-mode image does not contain enough information to determine the full shape of all three inclusions. During the segmentation process, the edges of the inclusions were accurate within the imaged region, but parts of each inclusion were cut off and assumed to be part of the background material. Meshes produced by manual segmentation are displayed in Fig. 6.

AutoP training was performed with the same data used for Phantoms 1 and 3 with known geometry. Cases (e) and (f) of Phantom 1 were trained in the same manner described in Section 2.4 with the exception of ANN sizes and a slight adjustment to pre-training. ANNs began with five nodes per hidden layer and a single node was added to each hidden layer in the fourth load step. Pre-training was performed using 1000 random stress-strain vectors instead of 100.

Case (e) of Phantom 3 required further adjustments to the ANN sizes and convergence criteria. For this training scenario, we initialized the ANNs with eight nodes per hidden layer and added a single node to each hidden layer in the sixth load step. Convergence criteria were set as $C_{max}^n=0.5$ and $C_{\mu}^n=0.15$ for load steps 1–3 and $C_{max}^n=0.2$ and $C_{\mu}^n=0.08$ for the remaining training steps.

The aforementioned changes to AutoP training were found through preliminary testing. With increased complexity in the stress and strain distributions caused by more complex shapes, it is likely the ANNs with larger hidden layers (i.e., larger capacity for information) were required to accurately describe the stress-strain relationship.

2.6 3-D Models with Known Interior Geometry

Our informational modeling approach was extended to 3-D models. For these tests, we returned to the assumption that the interior geometry was known. Data acquired on Phantom 4 was used to train ANN material models relating a six-component output stress vector (three normal components, three shear) to a six-component input strain vector.

Two data sets were acquired on Phantom 4. The first data set was obtained using the measurement setup described above (Fig. 1). A second set of compressions was performed where the US probe was rotated 90° about the x_2 axis. B-mode images from both data sets are displayed in Fig. 5d. Rotating only the probe allowed us to use the same FE mesh for the gelatin phantom and simply adjusting the orientation of the probe model. General 4-node, tetrahedral element (C3D4 in ABAQUS 6.13) were used to mesh the phantom, whereas the probe comprised 8-node, cubic elements (C3D8) and was modeled as a rigid body.

Measured forces were applied as distributed surface loads on the top surface of the probe in FEA_{σ} .

We trained ANN material models in separate instances of AutoP with the two different data sets. Once again, interior displacements were applied at three locations within the inclusion in FEA_{ϵ} . These points resided along the x_2 axis centered on the (x_1, x_3) plane and are highlighted in Fig. 5d. Displacements were applied to the same nodes for both data sets.

ANNs were pre-trained with 8000 random stress-strain vectors computed using a model of a 3-D Hookean, nearly incompressible ($\nu = 0.49$) solid. Each ANN started with eight nodes per hidden layer that was increased to nine in the sixth load step.

3 Results

3.1 2-D Phantoms with Known Geometry

Fig. 7 displays the reconstructed stress- and strain-vector maps for Phantom 1 case (d) computed in FEA_{ANN}^{10} . Images labeled axial, lateral, or shear strains (top row) and stresses (bottom row) are relative to the orientation of the ultrasound beam along the vertical x_2 axis. Again, the ANN material models were trained with AutoP using only surface force and displacement measurements and three internal displacement measurements within the inclusion.

Fig. 8 stress and strain images corresponding one to one with those in Fig. 7. The material properties used to compute these results were obtained from indentation measurements, so we refer to them as the FEA_{ind} results. They were generated using conventional FEA forward modeling under the same loading and boundary conditions as FEA_{ANN}^{10} for Phantom 1. The background and inclusion elements assume a linear-elastic material model with Poisson's ratio $\nu = 0.5$ and Young's modulus values $E = 9.16$ kPa and 22.9 kPa, respectively.

A direct comparison between strain estimates from conventional quasi-static elasticity imaging modalities can be made with Figs. 7b and 7d. The axial strain image shown in Fig. 7d was calculated using a least-squares strain estimator on the displacements computed from the speckle-tracking algorithm (Rivaz et al 2011). Conventional methods are limited to only the imaged region and are prone to artifacts, like the darker region around the inclusion.

Results shown in Figs. 7 and 8 are the outputs from FE analyses. The difference is that in Fig. 7 the element material properties were modeled by ANNs trained with AutoP using force-displacement data acquired on Phantom 1, whereas in Fig. 8 an ideal linear-elastic model was chosen using Young's modulus values estimated from indentation testing of Phantom 1 materials. Imposing a nonslip condition at the base of the phantom and a free-slip condition at the top surface, we were able to add stress and strain features in the image plane that should not and did not appear in the modulus image given a uniform background material. For example, the large variations in lateral strain (Fig. 7a) and stress (Fig. 7e) do not appear in the Young's modulus image. Note that the experimental conditions for the data in Fig. 7 are not exactly the same as the ideal conditions assumed in the modeled results of

Fig. 8. Therefore, we expect the corresponding pairs of stress-strain maps to be similar but not an exact match.

An essential test for trained ANN material models is their ability to accurately estimate stresses and strains in an object in response to an applied load. Comparing Figs. 7 and 8 we see modest differences between measured and modeled normal stresses σ_{11} , σ_{22} and normal strains ϵ_{11} , ϵ_{22} . Close agreement suggests that measuring axial forces and displacements at the surface and at a few points within the inclusions enabled ANNs trained with AutoP to learn to reliably predict normal stress and strains throughout a $5 \times 5 \text{ cm}^2$ area of a linear-elastic material with a spatial resolution on par with conventional FE modeling. *To the best of our knowledge, this is the first time that all relevant stress fields have been measured with this degree of accuracy and resolution without assuming the underlying mechanical properties, even for this relatively simple object.*

To compute the Young's modulus for these 2-D models, we used the incompressible result $\nu = 0.5$ and the equation

$$E(x_1, x_2) = \frac{(1 - \nu^2) \sigma_{22}(x_1, x_2)}{\nu \epsilon_{11}(x_1, x_2) + \epsilon_{22}(x_1, x_2)} \quad (5)$$

at each integration point the FE mesh, providing point-wise estimate of the Young's modulus, $E(x_1, x_2)$. In Eq. (5), $\sigma_{22}(x_1, x_2)$, $\epsilon_{11}(x_1, x_2)$, and $\epsilon_{22}(x_1, x_2)$ are the axial stress, lateral strain, and axial strain, respectively, at location (x_1, x_2) . This resulted in the modulus images labeled (h) Fig. 7 and (g) in Fig. 8. Because we supplied axial forces and displacements to AutoP during training and have incompressible materials, the axial and lateral stress and strain estimates are expected to be more accurate than shear estimates. Consequently, the second row of the plane-stress equation, which relates axial stresses to axial and lateral strains, was selected to estimate $E(x_1, x_2)$.

Table 2 lists the estimated values (from AutoP) and predicted values (from indentation) of Young's modulus for the four phantoms. The same table contains results where AutoP was trained using different orientations of the phantoms and different data sets for the known geometry cases. Additionally, Young's modulus estimates for 2-D models with estimated internal geometry and 3-D models with known geometry are included. Mean values measured from modulus images for Phantom 1–3 agreed with indentation estimates of the moduli within 20%. Fig. 9 are the Young's modulus images for Phantoms 2 and 3 computed using their respective stress and strain estimates produced in FEA_{ANN}^{10} . Described later is another method for Young's modulus estimation from trained ANNs that was used for the 3-D models.

In addition to evaluating the accuracy of the Young's modulus estimated from trained ANNs, we also wanted to ensure the deformations predicted by the ANNs in FEA_{ANN}^{10} were comparable to the measured displacements. Axial displacements at all nodes in the imaged region were compared those estimated via speckle-tracking. Errors reported in the last

column of Table 2 are $c_\mu \pm$ one standard deviation (Eq. 4) and c_{max} (Eq. 3) over all ten load increments. That is, displacements were estimated in FEA_{ANN}^{10} and c_μ and c_{max} were computed using all N nodes in the imaged region. A total of ten displacements were computed for each node and both errors were calculated using the $10*N$ values. The large standard deviation is attributed to larger displacement errors occurring during the first load increments where displacements are small. For example, in Phantom 1 case (a), c_μ calculated for FEA_{ANN}^{10} is 0.49 ± 0.04 using data from the first load increment only but decreases to 0.11 ± 0.05 when data from only the last load increment is considered. Similarly, c_{max} decreased from 0.58 in the first load increment to 0.22 at in the final increment.

3.1.1 Comparing Stress Response of Informational Models to a Constitutive

Matrix—We previously discussed and illustrated in Fig. 4 how an informational model essentially replaces the constitutive matrix to relate stresses and strains. The linear-elastic, isotropic material properties and geometry of the phantoms allow us to use the plane-stress approximation for the 2-D models. This means the response of the informational models developed for Phantoms 1–3 should describe the same behavior as the constitutive matrix in the plane-stress model.

ANN_{soft} for Phantom 1, case (d) was interrogated to obtain the stress response of the background material to various strains. The isolated strain vector $[e_{11} \ 0 \ 0]$ was used as input to ANN_{soft} . Values for e_{11} spanned the ± 0.05 range. After each vector was input the the ANN, the full stress response vector $[\sigma_{11} \ \sigma_{22} \ \sigma_{12}]$ was recorded. The same strain vectors were used in the plane-stress equation at the top of Fig. 10 to obtain the “ideal” stress response. In the plane-stress equation, the Young’s modulus was set to the value measured from macro-indentation (9.16 kPa) and $\nu = 0.5$. This process was repeated for isolated strain vectors $[0 \ e_{22} \ 0]$ and $[0 \ 0 \ e_{12}]$.

Stress responses computed from ANN_{soft} and the plane-stress model were plotted and arranged in Fig. 10 to mimic the 2-D constitutive matrix. In the plots, the dotted curves are the behavior described by ANN_{soft} and the solid curves are the response predicted by the plane-stress approximation. Perfect mechanical characterization by ANN_{soft} would result in the black curves falling on top of the red curves. We find agreement between the informational and constitutive models except for the shear plots of σ_{12} versus e_{12} , also evident when we compare Figs. 7c and 7g to Figs. 8c and 8f. We will return to this point later to discover how the shear response can be better learned by the informational models.

3.1.2 Varying the Amount and Uses of Training Data—Fig. 11 illustrates the influence of additional training data that convey material property information on ANN material model accuracy. All three plots show the axial-stress response σ_{22} to isolated lateral strains e_{11} that are input to different versions of ANN_{soft} for Phantom 1. Plot (a) are the response of ANN_{soft} trained with data from single set, case (a) in Table 1. The large disagreement between the solid- and dotted-curve data in Fig. 11a suggests not enough information was supplied during AutoP for ANN_{soft} to learn the relationship between axial stress and lateral strain. Adding independent training data from another view (cases (c) and

(d) in Table 1 greatly improve the ability of ANN_{soft} to characterize mechanical responses. Results of these AutoP trained models are shown in Figs. 11b and 11c.

We saw from the data in Figs. 7 and 10 that it is difficult to capture the shear stress response to shear strain by applying a uniaxial load and only measuring forces and displacements along the load axis. Although we do not apply a shearing stimulus, the phantom heterogeneity and external boundary conditions produce a small amount of shear, evidenced by the images of Figs. 8c and 8f.

Several versions of AutoP training were performed to evaluate their effects on the ability of ANN_{soft} to learn the shear response. These were: 1) training using axial force and displacement information (case (a) in Table 1; 2) providing both axial and lateral displacements during training; 3) imposing an assumption that the Jacobian matrix, $[\frac{\partial \sigma}{\partial \epsilon}]$, is symmetric during FE analysis; and 4) checking displacement errors at points not included in FEA_e during convergence testing. With the additional points in the displacement error calculation, a total of 65 points sampled uniformly across the imaged region were included in the convergence check for case (4).

While none of these techniques led to ANN_{soft} agreeing with the plane-stress model predictions, each of the adjusted training versions (2)–(4) improved the shear response compared to standard version (1). The greatest improvement occurred for version (4) when additional displacement data were included in the test for convergence. A comparison of the full shear stress field is shown in Fig. 12. Shown are the estimated shear stress distribution with no training adjustment (version(1)), enforcing stricter convergence testing (version (2)), and the behavior predicted by FEA_{ind} (Fig. 12c). The increased accuracy in Fig. 12b over Fig. 12a suggests there is merit in including displacement errors from points not supplied displacement data during FEA_e when performing the convergence check.

3.2 2-D Phantoms with Coarsely Estimated Geometry

Images of the Young's modulus distribution for cases (e) and (f) of Phantom 1 and case (e) of Phantom 3 are shown in Figs. 13a–13c, respectively. These images were computed in the same manner described before: the stress and strain fields computed in FEA_{ANN}^{10} were used in Eq.(5) to calculate the Young's modulus. Estimates of the Young's modulus of the background and inclusion materials for these three cases are included in Table 2.

Errors are present in the Young's modulus distributions because of geometry errors resulting from segmentation; however, the errors did not seem to propagate through training. That is, the informational models were still able to accurately characterize the gelatin materials even though geometry errors were present, evidenced by the close match in Young's moduli computed from the ANN material models and compared to those estimated from macro-indentation.

3.3 3-D Phantoms

Returning to the case where the interior geometry is known, the Young's modulus distribution in Phantom 4 with 0° probe rotation is shown in Fig. 13d. The modulus was

computed throughout the entire $50 \times 50 \times 50 \text{ mm}^3$ region; however, only two planes are shown for display purposes. All six components of stress and strain distributions were computed in FEA¹⁰_{ANN} and used in the following inversion based on the second row of the constitutive equation for a 3-D Hookean material:

$$E(x_1, x_2, x_3) = \frac{\sigma_{22}(1+\nu)(1-2\nu)}{\nu(\varepsilon_{11} - \varepsilon_{22} + \varepsilon_{33}) + \varepsilon_{22}} \quad (6)$$

where σ_{22} is the axial stress, $\nu = 0.5$, and ε_{11} , ε_{22} , and ε_{33} are the lateral, axial and elevational strains, respectively. For brevity in Eq.(6) the location (x_1, x_2, x_3) for the stress and strain components is implied.

Young's modulus values for Phantom 4 are reported in Table 2. Here, the reported Young's moduli were not computed as the mean of the regions in Fig. 13d. Instead, the ANN material models were probed directly to demonstrate alternative methods of estimating the mechanical parameters. Random values were chosen for all six strain components in the ± 0.01 range and used as input to the ANN material models. Strain values were randomized in order to avoid probing the ANNs with data used during training and the ± 0.01 range ensured all randomly generated strain vectors were within the bounds of the training data. The corresponding stress was computed and Eq.(6) solved to estimate the Young's modulus. A total of 300 random strain probes were input to each trained ANN to produce ten different estimates of the Young's modulus. The mean of these 300 values was taken as the estimate of the Young's modulus and included in Table 2.

4 Discussion

We have demonstrated the ability of informational models trained using the Autoprogressive method to learn the mechanical behavior of quasi-statically loaded, heterogeneous, linear-elastic, isotropic materials using a sparse sampling of force and displacement measurements. The results for the 2-D and 3-D phantom experiments show that all relevant stress and strain fields can be estimated with reasonable accuracy at a spatial resolution dictated by the FE mesh. The innovation in this approach is to replace the constitutive matrix normally used to solve the mathematical inverse problem with ANNs trained using the Autoprogressive method. The combination of finite-element methods and ANNs allows us to build models of mechanical properties without the need for simplifying assumptions.

Tests described in the Results section summarize our initial investigations at selecting measurement points for elasticity imaging. The procedure is to select as few points as possible to avoid the negative effects of noise in displacement estimation. We found that supplying displacement information at three points within each heterogeneity during FEA_e gave satisfactory results. Appendix 2 investigates the reason why only a few displacements are preferable when noise is present. These are combined with measurements of force and displacement at the surface where the ultrasound probe contacts the object. Such sparse sampling works well for quasi-static loading, where the applied force has time to propagate

throughout the medium. In that case the displacement at each point carries information about all points in the medium.

One ANN was trained for each material with distinct mechanical properties in the object, resulting in two informational models being developed for each phantom: ANN_{stiff} that characterized the inclusion(s) and ANN_{soft} for the background gelatin material. We found that incremental loading and diversification of the scanning views while measuring displacements at the same object locations improves informational model accuracy. Also, diversifying the locations of displacements estimated in the object for convergence testing allowed the ANN material models to learn the full range of stress and strain information more accurately. Interestingly, the training technique that worked well for Phantoms 1 and 2 having a single inclusion and meshed assuming the interior geometry was known, worked equally well on Phantom 3 with three inclusions, the 2-D models where geometry was estimated by manual segmentation, and the 3-D model of Phantom 4, albeit with minor adjustments to ANN sizes. Furthermore, errors in estimating the inclusion boundaries did not inhibit the ability of the informational models to learn the correct mechanical properties. Therefore it appears that greater object heterogeneity and geometry errors do not significantly increase the training burden. We will learn more about these types of issues in future studies involving objects with complex spatial heterogeneities and materials behaving in non-linear, path- and time-dependent manners.

One particularly exciting aspect of our technique is the ability to update existing informational models when new measurements are made. Coupling the ability of ANNs to learn new information with measurements taken at different times under different loading geometries, we can update an existing ANN with the new data instead of building a new model. Case (c) of Phantom 1 was an example of “retraining” previously developed informational models after new data was gathered. We found that providing new information to an existing informational model increased the accuracy of stress, strain, and Young’s modulus estimates. This capability allows for resampling of a material after the initial measurements if it is found that the informational models are not converging or to initiate training of models for new materials under investigation. However, care must be taken to ensure that the ANNs do not “forget” previously learned behaviors. Case (c) for Phantom 1 supplied a second set of force-displacement data for ANN retraining that was very similar to the first set of data. As such, there was not much risk of ANN_{soft} and ANN_{stiff} discarding the mechanical properties already learned. Potential problems associated with ANN retraining are most likely to occur when the new information is acquired under different loading situations or when the mechanical behaviors become complex. Fortunately, such issues were circumvented by interspersing previously learned data with the new data so that the ANN retain old information while learning new properties.

4.1 Effects of Limited Information

Stress and strain vectors produced during uniaxial compression of cubic phantoms mostly span the axial and lateral spaces. Limited shearing occurs, resulting in less information available for training the ANNs on the shear behavior. However, we showed it was possible for the ANNs to better learn the shear response by making adjustments to the training

process. The most successful adjustment was including more displacements in the error check for convergence. Forcing the ANNs to more accurately predict displacements not given during training creates a more challenging convergence criteria, which in turns increases the number of training iterations performed. It is true that this same scenario can be created by adjusting C_{max}^n and C_{μ}^n , but we then run the risk of overtraining the ANNs. Instead of forcing the ANNs to very accurately predict displacements at a few points, we relax the constraints and force the ANNs to approximate the displacements at many points.

4.2 Training Data Selection

There is much to be learned about what constitutes sufficient information to successfully build an informational model. The drawback of providing too much information when the measurement data are noisy was discussed previously. More interesting is that changing the nodes on which the model trains greatly affects the outcome of the informational modeling process. Three different cases illustrate this point, and in each case the only force measurements are those applied at the ultrasound probe-phantom interface.

Consider situations where displacements are given 1) at the probe only, 2) along a vertical line through the inclusion, and 3) along a horizontal line through the inclusion of Phantom 1. In all three cases, Young's modulus estimates of the background material were near the expected values; however, only in case (2) did we get accurate estimates for the Young's modulus of the inclusion. The result of case (1) is not completely unexpected since the background comprises the bulk of the phantom and will have the greatest effect on the force-displacement relationship measured by the US probe. Results of (2) and (3) start to reveal what type of information needs to be supplied for training.

It is evident that displacement information must be given for each region in the material with distinct mechanical properties. For the models in this study, probe displacements would be enough to characterize the axial stress-strain relationship of the gelatin comprising the background of the phantoms. Case (1) indicates that displacements must be provided for the inclusion. The same number of nodes were given displacement data in cases (2) and (3), the only difference lies in the orientation of the nodes compared to the loading direction. In (2), the given information lies along a line parallel to the load axis whereas in (3) the nodes given displacement data spanned a line perpendicular to the loading direction. Because only axial displacement data was given, the movement of the nodes in (3) reveals very little information about how the points move relative to each other after loading - indeed, if the given displacements contained no noise, all of the nodes would have very similar (or identical) axial motion. However, in (2), the nodes compress relative to each other, producing more information about the material properties of the inclusion.

We empirically find the best sampling strategy for these linear-elastic, isotropic materials is to measure displacements at nodes that vary significantly along the applied force gradient. For example, axial compression means the force gradient lies mostly along the x_2 axis. Therefore, we select nodes along the x_2 axis in which displacements vary to the greatest extent. It is important to note this sampling strategy may be insufficient for materials exhibiting more complex behavior. Future research of this method will include investigations

into the minimum requirements of force-displacement measurements to build informational models of materials that have nonlinear, path-dependent, time-dependent, and/or anisotropic mechanical properties.

4.3 Choice of Convergence Criteria

The choice of C_{μ}^n and C_{max}^n largely affect the number of iterations performed in each training step of AutoP. Enforcing very strict convergence criteria increases the risk of overtraining the ANNs, especially if the displacement data involved in the displacement error calculation contains high levels of noise. Conversely, relaxing the convergence criteria too much may decrease the accuracy of the ANN material models. At this point in time, the choice of convergence criteria is largely heuristic. But, the same criteria were used for all phantoms when the geometry was known, which may imply the choice for convergence criteria can be based on the complexity of the object. Case (e) of Phantom 3 - where the internal geometry was estimated via manual segmentation - was the only case where the convergence criteria changed. At the start, C_{μ}^n and C_{max}^n were relaxed, allowing for larger displacement errors. After training the ANNs over a few load increments, the both values were decreased, making convergence more difficult to achieve.

Gradually restricting the convergence criteria follows the idea that the ANNs more accurately model the mechanical behavior as training progresses. More iterations of AutoP means the ANNs are provided with an increasing amount of information. And, because the stress-strain estimates in FEA_{σ} and FEA_{ϵ} become more accurate as the ANNs learn the correct material properties, it is natural to enforce stricter requirements for convergence as the generated stress-strain training data becomes higher quality. In the case of Phantom 3, case (e), the estimated internal geometry produced more complex stress and strain distributions that required several iterations of AutoP for the ANNs to learn. But, as the ANNs better learned the mechanical properties, stronger requirements for convergence could be put in place.

Moving forward, the choice of C_{μ}^n and C_{max}^n likely depend on the geometric and material complexity of the object and accuracy of measured forces and displacements. Thus, convergence criteria will be determined empirically. Using AutoP on models of many different geometries and material property distributions may allow us to identify a trend.

4.4 Computational Load Considerations

With AutoP, most of the computation time is spent performing FEA_{σ} and FEA_{ϵ} . For two passes of ten load steps, a minimum of 60 FEAs are performed, up to 220 in the worst-case where the maximum number of iterations is reached during training for each load step. This assumes a minimum of one and a maximum of five training iterations per load step. Our current implementation of AutoP in Matlab 2013b required 51 ± 7 minutes to perform one pass of ten load steps in AutoP to create informational models of Phantoms 1–3 cases (a)–(d). However, after a few optimizations, AutoP was able to train 2-D models in ≈ 30 minutes. Conversely, developing 3-D models required ≈ 65 minutes. Note that these analyses were ran on a PC containing a quad-core processor operating at 3.4 GHz. To give an example of the time required to complete the FEAs, FEA_{ind} for Phantoms 1–3 (664 –

1146 DOF) was solved in ≈ 4 seconds whereas the solution time for Phantom 4 (7503 DOF) increased to ≈ 30 seconds.

There are many ways to reduce the run time of AutoP, including the use of dedicated hardware and software and through optimized sampling and training strategies. Of particular interest is offloading computation to a GPU, which could provide up to 100 $\times$ increase in speed. Other groups have been successful in developing real-time FEA solvers for biomedical applications (Taylor et al (2007)) and structural analysis (Huang et al (2015)) with the use of a GPU. It is our belief that it will be possible to develop informational models in real-time during imaging. Methods to decrease the training time with AutoP will be investigated after we make significant progress in the basic research of this method.

5 Conclusions

Informational models based on AutoP have an ability to learn quasi-static mechanical behavior from measurements of surface forces and sparse surface and internal displacement information. This approach has distinct advantages when the mechanical properties of the medium are not well characterized, which is the situation in medical elasticity imaging. We trade the conventional parametric inverse problem approach involving numerous assumptions for a nonparametric machine-learning technique. This technique learns to model stress and strain fields from which imaging parameters are found.

The ability of informational models to accurately predict mechanical behavior is limited by the type and amount of force-displacement information provided. Model accuracy can be improved by adding data from different spatial views and carefully distributing measurement data for FEA training and convergence testing. An advantage that AutoP offers in studying similar objects is its ability to accumulate knowledge, which minimizes training time in ways we are just beginning to understand. We believe this novel computational approach has much to offer for imaging mechanical properties of biological media.

Acknowledgments

Research reported in this publication was supported by the National Cancer Institute of the National Institutes of Health under Award Number R01CA168575. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

The authors gratefully acknowledge a generous supply of gelatin provided by Rousselot Inc., Dubuque IA.

The authors would also like to thank H. Rivaz, E. M. Boctor, M. A Choti, and G. D. Hager for providing the source code for their real-time strain estimation algorithm.

References

- Altahhan KN, Wang Y, Sobh N, Insana MF. Indentation measurements to validate dynamic elasticity imaging methods. *Ultrasonic Imaging*. 2016; 38(5):332–345. [PubMed: 26376923]
- Barr RG. Elastography in clinical practice. *Radiologic Clinics of North America*. 2014; 52(6):1145–1162. [PubMed: 25444097]
- Ghaboussi, J. *Advances of Soft Computing in Engineering*. Springer; 2010. Advances in neural networks in computational mechanics and engineering; p. 191-236.

- Ghaboussi J, Sidarta D. New nested adaptive neural networks (NANN) for constitutive modeling. *Computers and Geotechnics*. 1998; 22(1):29–52.
- Ghaboussi J, Pecknold DA, Zhang M, Haj-ali RM. Autoprogressive Training of Neural Network Constitutive Models. *International Journal for Numerical Methods in Engineering*. 1998; 42(1):105–126.
- Hashash YMA, Jung S, Ghaboussi J. Numerical implementation of a neural network based material model in finite element analysis. *International Journal for Numerical Methods in Engineering*. 2004; 59(7):989–1005.
- Hoerig C., Ghaboussi, J., Insana, MF. Informational modeling of tissue-like materials using ultrasound. 12th IEEE International Symposium on Biomedical Imaging, ISBI 2015; Brooklyn, NY, USA. April 16–19, 2015; 2015. p. 239-242.
- Hoerig C, Ghaboussi J, Fatemi M, Insana MF. A new approach to ultrasonic elasticity imaging. *Proc SPIE Medical Imaging: Ultrasonic Imaging and Tomography*, International Society for Optics and Photonics. 2016; 9790:G1–9.
- Huang J, Ong S, Nee A. Real-time finite element structural analysis in augmented reality. *Advances in Engineering Software*. 2015; 87:43–56.
- Kallel F, Bertrand M. Tissue elasticity reconstruction using linear perturbation method. *IEEE Transactions on Medical Imaging*. 1996; 15(3):299–313. [PubMed: 18215911]
- Orescanin M, Insana M. Shear modulus estimation with vibrating needle stimulation. *Ultrasonics, Ferroelectrics, and Frequency Control*, IEEE Transactions on. 2010; 57(6):1358–1367.
- Riedmiller, M., Braun, H. A direct adaptive method for faster backpropagation learning: the RPROP algorithm. *IEEE International Conference on Neural Networks*; IEEE; 1993. p. 586-591.
- Rivaz H, Bector EM, Choti MA, Hager GD. Real-time regularized ultrasound elastography. *IEEE Transactions on Medical Imaging*. 2011; 30(4):928–45. [PubMed: 21075717]
- Taylor, ZA., Cheng, M., Ourselin, S. *Real-Time Nonlinear Finite Element Analysis for Surgical Simulation Using Graphics Processing Units*. Springer Berlin Heidelberg; Berlin, Heidelberg: 2007. p. 701-708.
- Tyagi M, Goenezen S, Barbone PE, Oberai AA. Algorithms for quantitative quasi-static elasticity imaging using force data. *International Journal for Numerical Methods in Biomedical Engineering*. 2014; 30(12):1421–1436. [PubMed: 25073623]

6 Appendix 1: Accuracy of displacement estimates

Building accurate informational models depends on acquiring accurate force and displacement measurements. We were interested in evaluating the displacement errors introduced through speckle-tracking and how those errors propagated through the load increments. Because displacement estimates were made between consecutive RF frames, estimates for the later load increments are the sum of many displacement maps. Quantifying displacement errors and tracking how they changed with successive load increments allowed us to ensure the effects of displacement errors on ANN training with AutoP would be minimal.

A $50 \times 50 \times 50 \text{ mm}^3$ homogeneous phantom was constructed and imaged in the same manner described in Section 2.1. After creating the FE model (here, a rectilinear mesh was used for the phantom), a version of FEA_e was solved for the model where only probe displacements were applied. A total of 10 load increments were applied to compress the phantom 1.5mm. Because the phantom was homogeneous and the FEA was displacement controlled, the displacements computed at the nodes are the same regardless of the stiffness of the phantom.

Axial node displacements estimated from the US images were compared to those computed in the FEA using the equation

$$e_i^n = \frac{|u_i^n - \hat{u}_i^n|}{|max(u_i^n)|} \quad (7)$$

where i indexes the nodes, n indexes the load increment, e_i^n is the displacement error, u_i^n is the displacement computed in the FEA, and $\hat{u}_i^n$ is the displacement estimated via speckle tracking. This is the same error computation for convergence described in Section 2.2.4. A histogram of the axial displacement errors for all 399 nodes in the imaged region across the ten load increments is plotted in Fig. 14a. To compare these errors to the convergence criteria specified in Sections 2.4 and 2.6, $c_\mu = 0.000022$ and $c_{max} = 0.0001$. These values are several orders of magnitude smaller than the chosen convergence criteria, suggesting that the test for convergence in each iteration of AutoP will be mostly affected by the material properties learned by the ANNs, not errors caused by speckle-tracking.

Tracking c_μ and c_{max} through the load increments, it was found that both increased monotonically and approximately linearly, as shown in Fig. 14b. A monotonically increasing error implies that the displacement errors are biased and the linearity suggests the error in displacement estimates between RF frames is consistent. But, as just mentioned, the errors are well below the set convergence criteria and should not have a significant impact on the displacement error calculation. It is a point to consider, though, if applied loads produce large deformations over many increments.

7 Appendix 2: The effect of noise on stress-strain training data

Classic mathematical inverse methods typically require many samples of data to reduce the effect of noise on the solution. It seems counter-intuitive, then, when we claim that using a smaller number of displacements in FEA_e enables the ANNs to more reliably learn the underlying mechanical properties. Our new method of elasticity imaging takes an alternative approach to the inverse problem that does not appear to follow the rules of classical inverse methods.

Current elastographic modalities formulate inverse problems that involve first or second order derivatives that are highly sensitive to noise. To combat noise amplification, those methods use sums or averages over many samples. We, however, are not directly solving the inverse problem. The sparse measurements of force and displacements are distributed across the entire mesh of the two FEA elements in AutoP and iteratively reconciled through the ANNs. Since AutoP model building occurs after several load steps and passes, the influence of noise introduced by sparse measurements at each load step is diminished through the process of FEA reconciliation. The effect of noise in the pipeline can best be demonstrated with a small example.

Using measurements from Phantom 1, consider two situations: 1) axial displacements are provided at every node in the imaged region (Fig. 15a, left) and 2) axial displacements are provided at three locations within the inclusion (Fig. 15a, right) in FEA_e. The two axial strain images in the center of Fig. 15a visualize the effect of noise on the resulting strain.

Note that the color scales are different to help visualize the axial strain distribution in both cases. Going one step further, the axial stress-strain data is plotted in Fig. 15b for ANN_{soft} and ANN_{stiff} . These plots were obtained by pairing stresses from FEA_{σ} with strains from FEA_{ϵ} and pulling out only the axial components.

It is the job of the ANNs to generalize the data in Fig. 15b. Note that the stress and strain vectors each have three components in 2-D (six in 3-D), meaning the Fig. 15b only provides a partial picture of the function being approximated by the ANNs. However, the plots do provide a useful comparison. In the case of many displacements, the increased variance present in the strain data causes the ANNs to learn incorrect material properties. For the gelatin phantoms exhibiting linear-elastic material properties, the spreading of the strain data decreases the slope of the stress-strain curve. In all of our tests, using displacements at every point (or every other point) in the imaged region resulted in Young's modulus estimates much smaller than the true value. Using fewer displacements produces *some* variance that influences changes in the ANN connection weights during training. As AutoP iterations progress, the smaller variations caused by using few displacements in FEA_{ϵ} allow the ANNs to gradually learn the correct material properties instead of attempting to force large changes in the learned behavior at once.

Fig. 1.

Diagram of the experimental setup using a single-inclusion gelatin phantom scanned along the cross section of the inclusion. A centrally-positioned ultrasound probe coupled to a force/torque transducer (not shown) is maneuvered by a 3-D positioning system. The phantom is placed on a fixed non-slip base and is compressed from above with surface force p as an RF echo frame (dashed line region) is acquired. Speckle-tracking applied to echo signals measures displacements within the RF echo frame. Locations marked 'x' indicate positions where displacements were applied in FEA_e for training. These displacements are also used for convergence testing. Locations marked 'o' are examples of locations where additional displacement data can be used for FEA_e and/or convergence testing.

Fig. 2.

Flow diagram of the Autoprogressive method. Measured forces and displacements from one load increment are supplied to finite element analyses FEA_σ and FEA_ϵ . In the i^{th} iteration, current version of the ANNs, $[NN]_i$, are used as the element material models to solve the FEAs, producing estimates of the stress (σ_i) and strain (ϵ_i) vectors throughout the model. Stress-strain pairs are collected and used to update the connection weights of the ANNs, producing the updated network $[NN]_{i+1}$. This process iterates, and each time displacement estimates $\hat{u}_i$ from FEA_σ are compared with measured displacements u_i to compute error vectors Δu_i . If displacement errors are above set thresholds, the training iterations continue using the same force-displacement data. Once errors fall below the thresholds, training for the i^{th} iteration has converged. This training process is performed for all force-displacement data sets to complete the first pass and may be repeated for the same or new data in a second pass.

Fig. 3.

Example of the evolution of the mechanical behavior described by an ANN during training with AutoP. After pre-training (-...-), the stress-strain response of the ANN is significantly different than the true mechanical behavior (—). Performing training iterations using force-displacement data from the first load increment (Load Step 1) gradually moves the stress-strain curve produced by the ANN toward the true solution. After completing training iterations using the data from the first load increment, the ANN accurately describes the mechanics only over the strain range produced for the first step (- - -). The valid range of the ANN is increased by continuing training through load step 2 (- · - ·).

Fig. 4.

After training with AutoP, an ANN material model is able to compute an output stress vector in response to an input strain, essentially replacing the constitutive matrix. The same information described by the components of the constitutive matrix is contained non-parametrically in the connection weights of the ANN. Using the ANN in a computational experimental like a FEA allows the full stress and strain vectors to be estimated throughout an entire material in response to an externally applied force. Access to estimates of all components of stress and strain allows mechanical parameters from any constitutive model to be estimated.

Fig. 5.

FE meshes overlaid on B-mode images of (a) Phantom 1, (b) Phantom 2, and (c) Phantom 3. The highlighted nodes in (a)–(c) indicate locations where displacement information was provided during training with AutoP. (d) Orientations of the B-mode images in the 3-D model. For developing 3-D ANN material models, we once again used only three internal displacements. Displacements were provided along the line of intersection of the two B-mode images and given at the same points for both US probe orientations.

Fig. 6.

FE meshes produced by manually segmenting B-mode images. (a) Segmentation of Phantom 1 closely matching the internal geometry (case (e)). The RMS error between the true inclusion boundary and the estimated boundary was 0.29mm. (b) Poor segmentation of Phantom 1 where the inclusion boundaries significantly differ from the true boundaries (case (f), RMSE = 0.76mm). The thick, black line indicates the boundary of the inclusion in the mesh. (c) Segmentation of Phantom 3 (case (e), RMSE = 0.76mm). The B-mode image does not capture the full shape of the inclusions. The boundaries of the inclusion within the B-mode are accurately captured, but inclusion boundaries outside of the imaged region are cut off.

Fig. 7.

Stress and strain fields computed in FEA_{ANN}^{10} for Phantom 1, case (d). Trained ANNs were used as the element material models to solve a forward FEA problem to produce estimates of the full stress and strain vectors. The results displayed were interpolated to a uniform grid. (a) Lateral strain. (b) Axial strain. (c) Shear strain. (d) Axial strain computed from speckle-tracking applied to the acquired RF frames. (e) Lateral stress. (f) Axial stress. (g) Shear stress. (g) Young's modulus distribution $E(x_1, x_2)$ computed using Equation 5.

Fig. 8.

FEA_{ind} estimates of the stress and strain fields for Phantom 1 under plane-stress (CPS4 elements in ABAQUS) using Young's modulus values estimated from macro-indentation tests. A linear-elastic material model replaced each ANN for the mesh in Fig. 7 to solve FEA_{ind} and produce these baseline estimates. These are the forward FEA estimates that we use to compare with the AutoP trained models shown in Fig. 7. (a) Lateral strain. (b) Axial strain. (c) Shear strain. (d) Lateral stress. (e) Axial stress. (f) Shear stress. (g) Young's modulus image $E(x_1, x_2)$ computed using Equation 5.

Fig. 9. Young's modulus images computed using Equation 5 with stresses and strains estimated in FEA_{ANN}^{10} for (a) Phantom 2 and (b) Phantom 3.

Fig. 10.

Stress response of a trained ANN_{soft} material model to the strain probes $[\varepsilon_{11} \ 0 \ 0]$, $[0 \ \varepsilon_{22} \ 0]$, and $[0 \ 0 \ \varepsilon_{12}]$. For each vector, the value of the strain component swept the ± 0.05 range and the full stress vector was recorded. Slopes of the stress-strain curves should match the corresponding component of the constitutive matrix for a plane-stress model having a Young's modulus of 9.16 kPa (estimated from macro-indentation) and assumed Poisson's ratio of 0.5. Dotted curves correspond to the stress response of the ANN whereas the solid curves were computed by applying the strain vectors to the plane-stress model.

Fig. 11.

Comparison of the σ_{22} response to an isolated lateral strain (ϵ_{11}) input. The dotted curves are the recorded response of the ANN modeling the background gelatin material for Phantom 1 (a) case (a); (b) case(c); and (c) case (d). As with Figure 10, the solid curves were computed using a plane-stress approximation.

Fig. 12.

Shear stress fields (σ_{12}) computed in FEA_{ANN} after training versions (1) and (4) for Phantom 1. (a) Training version (1) corresponds to Phantom 1, case (a) in Table 1. (b) Training version (4) where displacement convergence was required for additional nodes not provided in FEA_e ; (c) Shear stress computed for an ideal linear-elastic plane-stress model using a Poisson's ratio of 0.5 and Young's moduli estimated from macro-indentation.

Fig. 13.

Young's modulus images computed after training on (a)–(c) 2-D models where segmentation was used to identify internal boundaries and (d) a 3-D model. (a) Good segmentation of the internal structure of Phantom 1. (b) Poor segmentation of the internal structure of Phantom 1. (c) Segmented version of Phantom 3. (d) Two planes of the Young's modulus distribution calculated from FEA_{ANN}^{10} of Phantom 4 case (a).

Fig. 14.

(a) Histogram of displacement errors. (b) Both c_μ^n and c_{max}^n , where n is the load increment, increase approximately linearly, suggesting the speckle-tracking produces a consistent, biased error.

Fig. 15.

Comparing the effects of using many noisy displacements to only a few. Using many noisy displacements (a) causes large variance in the stress-strain training data (b) that have a negative impact on the ability of the ANNs to learn the underlying mechanical properties.

Table 1

Six different training scenarios for AutoP. These are the cases referred to in Table 2 for Phantoms 1 and 3 (Fig. 5). Data set 1 are measurements for a centrally-located US probe (Fig. 1). Data set 2 are measurements for an off-center US probe location.

Cases for Phantom 1, Known Geometry	
Case (a)	training using data set 1 only
Case (b)	training using data set 2 only
Case (c)	training using data set 1 and then retraining on set 2
Case (d)	training using both sets simultaneously
Cases for Phantoms 1 and 3, Estimated Geometry	
Case (e)	Good segmentation - boundaries closely followed
Case (f)	Poor segmentation - large errors in boundary estimation

Biomech Model Mechanobiol. Author manuscript; available in PMC 2018 June 01.

Young's Modulus Measurements of Phantom Components					Displacement Errors c_p/c_{max}
{Phan. #} (Case) [Rot. Angle]	ANN Estimated (kPa)		Indentation Meas. (kPa)		
	Background	Inclusion	Back.	Inc.	
2-D Models, Known Geometry					
{1} (a) [0 deg]	8.01 ± .5599	19.35 ± .2022			0.17 ± 0.13 / 0.58
{1} (b) [0 deg]	11.01 ± .1189	27.71 ± .5141	9.16	22.9	0.12 ± 0.10 / 0.48
{1} (c) [0 deg]	10.18 ± .3408	22.15 ± .3711			0.16 ± 0.11 / 0.55
{1} (d) [0 deg]	9.06 ± .1433	22.39 ± .3883			0.19 ± 0.14 / 0.65
{2} (a) [0 deg]	8.57 ± .5596	28.84 ± .1770	8.95	26.87	0.17 ± 0.10 / 0.42
{2} (a) [270 deg]	9.04 ± .0561	26.19 ± .1315			0.12 ± 0.11 / 0.45
{3} (a) [0 deg]	7.99 ± .2755	26.65 ± .2105			0.19 ± 0.07 / 0.35
{3} (a) [90 deg]	9.62 ± .1132	25.58 ± .3166	8.00	24.58	0.11 ± 0.12 / 0.68
2-D Models, Estimated Geometry					
{1} (e) [0 deg]	9.01 ± 0.10	26.65 ± 0.21	9.16	22.9	0.17 ± 0.13 / 0.70
{1} (f) [0 deg]	9.18 ± 0.10	22.42 ± 2.28			0.16 ± 0.11 / 0.63
{3} (e) [0 deg]	10.98 ± 0.47	24.24 ± 0.38	8.00	24.58	0.22 ± 0.17 / 0.81
3-D Models, Known Geometry					
{4} (a) [0 deg]	11.02 ± 0.01	25.22 ± 0.05			0.10 ± 0.10 / 0.51
{4} (a) [90 deg]	11.02 ± 0.01	25.23 ± 0.06	11.37	23.25	0.14 ± 0.12 / 0.46