PAPER • OPEN ACCESS

Ensemble Kalman inversion for magnetic resonance elastography

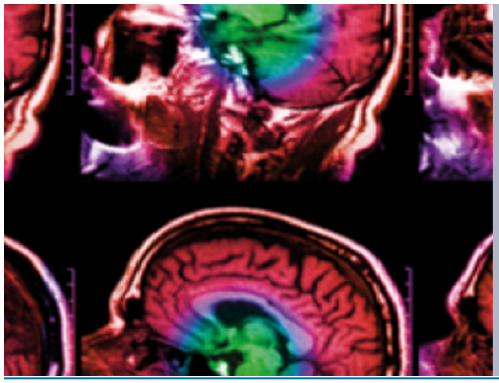

IOP Publishing Phys. Med. Biol. 67 (2022) 235003 https://doi.org/10.1088/1361-6560/ac9fa1

Physics in Medicine & Biology

IPEM Institute of Physics and Engineering in Medicine

OPEN ACCESS

RECEIVED
8 April 2022

REVISED
14 October 2022

ACCEPTED FOR PUBLICATION
1 November 2022

PUBLISHED
25 November 2022

PAPER

Ensemble Kalman inversion for magnetic resonance elastography

Marco Iglesias[1], **Deirdre M McGrath**[2,3], **M V Tretyakov**[1] **and Susan T Francis**[2,3]

[1] School of Mathematical Sciences, University of Nottingham, Nottingham, United Kingdom
[2] NIHR Nottingham Biomedical Research Centre, Nottingham University Hospitals NHS Trust and University of Nottingham, Nottingham, United Kingdom
[3] Sir Peter Mansfield Imaging Centre, University of Nottingham, Nottingham, United Kingdom

E-mail: marco.iglesias@nottingham.ac.uk

Keywords: magnetic resonance elastography, ensemble Kalman inversion, Bayesian inversion

Abstract

Magnetic resonance elastography (MRE) is an MRI-based diagnostic method for measuring mechanical properties of biological tissues. MRE measurements are processed by an inversion algorithm to produce a map of the biomechanical properties. In this paper a new and powerful method (ensemble Kalman inversion with level sets (EKI)) of MRE inversion is proposed and tested. The method has critical advantages: material property variation at disease boundaries can be accurately identified, and uncertainty of the reconstructed material properties can be evaluated by consequence of the probabilistic nature of the method. EKI is tested in 2D and 3D experiments with synthetic MRE data of the human kidney. It is demonstrated that the proposed inversion method is accurate and fast.

1. Introduction

Magnetic resonance elastography (MRE) McGrath (2018) is an MRI-based diagnostic method used to measure mechanical properties of biological tissues. Changes in these properties, in particular elasticity, can be associated with disease such as cancer (Sinkus *et al* 2000, Venkatesh *et al* 2008, Li *et al* 2011, Sakai *et al* 2016) and fibrosis (Yin *et al* 2007, Singh *et al* 2015). The two main classes of MRE are dynamic or harmonic (Muthupillai *et al* 1995), and static or 'quasi-static' methods (McGrath *et al* 2012). In dynamic MRE mechanical waves are produced in the tissue using a vibrating driver, while in static MRE the whole tissue volume is compressed. In each case, the resulting displacement field is measured with motion encoding gradients used to detect the displacements as phase-shifts in the MRI images received by the scanner. These images are then processed by an inversion algorithm to produce a map of the biomechanical properties of the tissue, typically referred to as an *elastogram*. Dynamic MRE is now being widely used for clinical applications in the liver (Singh *et al* 2015, Mathew and Venkatesh 2018), and there is growing interest in applying the method in a wide range of organs including the spleen and kidney (Venkatesh *et al* 2013), brain (Hiscox *et al* 2016), lung (Marinelli *et al* 2017), and muscle (Kennedy *et al* 2017).

The main elasticity parameter of interest in dynamic MRE is the shear elastic modulus, which varies by several orders of magnitude across different tissues and between the healthy and diseased state. Most soft tissues have material properties that are intermediate between those of solids and fluids, and can be characterised by viscoelastic constitutive material models. The complex shear modulus G summarises the shear viscoelasticity: the real component, or storage modulus G_s, defines the solid-like shear elasticity and the imaginary component, or loss modulus G_l, defines the fluid-like shear viscosity. In this context, an inversion algorithm is employed to infer G_s and G_l from MRI measurements of tissue displacements.

A variety of inversion algorithms have been explored for MRE (see e.g. Van Houten *et al* 1999, Papazoglou *et al* 2008, Doyley 2012, McGarry *et al* 2013, Ammari *et al* 2015, Fovargue *et al* 2018, McGarry *et al* 2019 and references therein). One of the more widely used methods is direct inversion (DI), or algebraic inversion (Papazoglou *et al* 2008). Other methods include local frequency estimation and iterative based methods (Manduca *et al* 1996, Van Houten *et al* 1999, Manduca *et al* 2001, Hu and Shan 2020). The quality and accuracy

of calculated elastograms is limited by the assumptions and mathematical methods employed. For example, in direct or algebraic inversion, typically the assumption of local tissue homogeneity is made (Sinkus *et al* 2005), and this can lead to large errors in the elastograms at the boundaries between tissues with very different underlying elastic properties. Also, algebraic inversion involves numerical differentiation of imaging data, which can amplify imaging-related noise. Iterative methods, on the other hand, avoid differentiation of data by posing the inversion as a least-squares problem that is solved numerically. Since this problem is inherently unstable (or ill-posed), regularisation strategies are needed in order to compute a stable reconstruction of the stiffness properties. However, most existing regularisation strategies produce over-smoothed properties which cannot accurately reconstruct the interface between healthy tissue and malignancies. While iterative methods can be equipped to address the instabilities (or ill-posedness) of the inversion in MRE, these methods are computationally more costly than those used in DI. Consequently, a large body of work has been devoted to improving the computational efficiency of iterative methods by using a domain decomposition (subzone) method (Van Houten *et al* 1999, McGarry *et al* 2013, 2019).

Most existing approaches, either direct or iterative, for inversion in MRE are developed within deterministic frameworks which do not account for the uncertainty in the produced elastograms. This uncertainty arises from the unavoidable presence of measurement noise as well as the fact that the solution to the inversion problem in MRE may not be unique. In other words, multiple elastograms may be consistent with the same measured data. While rigorous theoretical work on uniqueness in MRE reconstructions is very limited (Higashimori 2007), existing work has proven uniqueness under idealised mathematical hypotheses that are not likely to be satisfied in relevant settings for MRE (i.e. for stiffness properties in the presence of malign tissues).

In this work we propose a new inversion framework for MRE via the ensemble Kalman inversion (EKI) algorithm developed in Iglesias *et al* (2013), Iglesias (2016), Iglesias and Yang (2021), and that has been successfully used for tomography and detection in various applications including geophysics (Tso *et al* 2021) and engineering (Simon *et al* 2018, Iglesias *et al* 2018, Matveev *et al* 2021). The proposed EKI algorithm for MRE is derived from the Bayesian approach for inverse problems (Kaipio and Somersalo 2005, Stuart 2010). Therefore, contrary to most existing inversion algorithms for MRE, where a single estimate is produced with no known measure of its uncertainty, the proposed EKI-based framework provides a full characterisation of the posterior probability distribution of the inferred tissue properties. From this distribution we are able to compute measures of the uncertainty of the storage and loss modulus such as the (pointwise) variance and credible intervals as well as the probability for the presence of anomalies within the tissue under investigation.

In contrast to DI approaches which rely on the assumption of local homogeneity (i.e. piecewise constant stiffness properties) or standard iterative methods which often produce over-smoothed elastograms, our approach characterises the storage and loss modulus G_s and G_l via spatially-varying (fully heterogeneous) random functions that may display abrupt changes in the presence of unknown anomalies (e.g. tumours or cysts). We allow for these potential abrupt changes by encoding a level-set parameterisation (Iglesias *et al* 2016, Dunlop *et al* 2017, Chada *et al* 2018) within the EKI algorithm which enables us to statistically identify well-defined boundaries between normal and anomalous tissue.

Here we implement and test the proposed EKI-based MRE inversion framework using an underlying forward model based on a standard isotropic linear viscoelastic model (Fovargue *et al* 2018, Sinkus *et al* 2005, Li *et al* 2021) that we program in MATLAB via the PDE toolbox `pdetool`. Nonetheless, since one of the advantages of the EKI framework is that the forward model is used in a black-box fashion, the proposed EKI-based algorithm can be easily used with other acoustic wave propagation solvers.

We illustrate capabilities of the proposed EKI-based MRE approach via synthetic (virtual) MRE experiments of the kidney in 2D and 3D. We demonstrate that the proposed EKI-algorithm for MRE produces reliable probabilistic estimates of the stiffness properties of the tissue (including sharp discontinuities from the presence of disease) together with measures of uncertainty that can provide practitioners with a valuable tool for decision making under the presence of inherent uncertainties in the stiffness properties of biological tissue.

In section 2 we introduce a standard formulation for harmonic MRE—a linear viscoelastic PDE problem which is the forward model in this paper. We then we formulate the MRE inverse problem and briefly review some of the existing methods for solving this, including an explanation of the Bayesian approach to inverse problems which we exploit in this paper. The key for success of the Bayesian approach lies in an appropriate parameterisation of the quantities of interest (i.e. stiffness properties of the tissue) which is proposed in section 3. We illustrate the proposed Bayesian approach on 2D and 3D experiments in section 4. Concluding remarks are in section 5. Appendices provide technical details about the parametrisation and the prior used. *Supplementary material* contains additional experiments demonstrating how the accuracy of the proposed approach depends on (i) choice of the prior ensemble (including ensemble size), (ii) level of measurement noise, and (iii) segmentation uncertainties.

2. Bayesian reconstruction of stiffness properties

In this section we formulate the forward and inverse problems for MRE (section 2.1). We discuss some of the existing inversion approaches which have been used for MRE (section 2.2), and outline the Bayesian approach (section 2.3).

2.1. Forward and inverse problems for MRE

This section introduces the standard formulations for harmonic MRE, we refer the reader to Li *et al* (2021), Sinkus *et al* (2005), Jiang *et al* (2011), Ammari *et al* (2015), Fovargue *et al* (2018) (and references therein) for further details. Our aim is to describe deformation of a tissue occupying a physical domain $\Omega \subset \mathbb{R}^d$ ($d = 2, 3$) via excitation of a harmonic wave with angular frequency ω. The tissue is assumed to be isotropic and to have density ρ. A linear viscoelastic model is employed in which the amplitude of the tissue displacement, denoted by $\mathbf{u}(\mathbf{x})$, $\mathbf{x} \in \Omega$, is governed by the following PDE (see e.g. Landau and Lifschitz 1986, Ciarlet 1988, Jiang *et al* 2011, Ammari *et al* 2015, McGrath 2018, Jiang *et al* 2020 and references therein):

$$\rho\omega^2\mathbf{u} + \nabla \cdot [G(\nabla\mathbf{u} + (\nabla\mathbf{u})^T)] + \nabla(\lambda\nabla \cdot \mathbf{u}) = \mathbf{0}, \qquad \text{in } \Omega, \tag{1}$$

where $\lambda(\mathbf{x})$ and $G(\mathbf{x})$ are the (complex) first and second Lame parameters, respectively. Equation (1) has to be accompanied by some boundary conditions (see e.g. Ciarlet 1988, Jiang *et al* 2011, Ammari *et al* 2015, Jiang *et al* 2020 and also section 4.1 here).

For the MRE setting we note that ω is a known/controlled variable, while tissue density, ρ, is assumed to be that of water (i.e $\rho \approx 10^3$ kg m^{-3}). Once we specify the tissue stiffness properties characterised via $\mathbf{S}(\mathbf{x}) \equiv (G(\mathbf{x}), \lambda(\mathbf{x}))$, the forward problem consists of solving equation (1) supplemented with appropriate boundary conditions. We can formulate the forward problem in terms of the so-called *forward map* denoted by $\mathcal{F}$:

$$\mathbf{S} = (G, \lambda) \Longrightarrow \mathcal{F}(\mathbf{S}) = \mathbf{u}_M \equiv (\mathbf{u}(\mathbf{x}_1), \ \ldots, \ \mathbf{u}(\mathbf{x}_M)), \tag{2}$$

which for every admissible pair (G, λ) of stiffness properties returns the solution to (1) evaluated at M points, denoted by $X_M \equiv \{\mathbf{x}_m\}_{m=1}^M$. We define these points so that we can subsequently confront our modeling solutions with observed data collected on a given mesh/grid. In expression (2) we have removed the explicit dependence of $\mathbf{S} = (G, \lambda)$ on $\mathbf{x}$ to emphasize that it is the forward map. Furthermore, note that $\mathcal{F}$ is a nonlinear operator that, in general, cannot be expressed in a closed form.

With the aid of the forward map, we pose the following inverse problem for MRE: given a complex vector, $\mathbf{d}^\dagger$, of noisy MRI measurements of tissue deformation (measured on X_M) find tissue stiffness properties $\mathbf{S}^\dagger = (G^\dagger, \lambda^\dagger)$. A key underlying assumption in formulating this inverse problem is that $\mathbf{d}^\dagger$ are measurements of displacements $\mathbf{u}_M^\dagger$ simulated according to (1) using the truth (i.e. $\mathbf{S}^\dagger = (G^\dagger, \lambda^\dagger)$). One often assumes, for example, that

$$\mathbf{d}^\dagger = \mathbf{u}_M^\dagger + \eta^\dagger = \mathcal{F}(\mathbf{S}^\dagger) + \eta^\dagger, \tag{3}$$

where $\eta^\dagger$ is an unknown measurement error/noise. One fundamental property of this inverse problem is that it is ill-posed in the sense of continuity (Engl *et al* 1996, Kaipio and Somersalo 2005). In other words, one can find a sequence of stiffness properties $\{\mathbf{S}_n\}_{n=1}^\infty$ so that the sequence $\mathcal{F}(\mathbf{S}_n)$ converges to $\mathcal{F}(\mathbf{S}^\dagger)$ while $\mathbf{S}_n$ does not converge to $\mathbf{S}^\dagger$. Furthermore, the solution to the inverse problem may not be unique and, hence, multiple solutions (stiffness properties) may lead to the same observed data. Even in the idealised case where the measurement noise is absent, existing theoretical work ensures uniqueness only under very specific and often unrealistic conditions (e.g. for smooth stiffness properties Higashimori 2007). Thus, there is a need to quantify uncertainty of the reconstructed stiffness properties that we produce via solving the inverse problem.

2.2. Existing inversion approaches

Direct inversion (DI) is one of the most commonly used inversion methodologies for MRE (see e.g. Sinkus *et al* 2005, Papazoglou *et al* 2008). The following outlines the main idea of DI and notes that that there are many variants of DI including those that use weak formulations of (1).

In the context of the forward map defined earlier, DI neglects measurement errors and seeks a direct solution of the nonlinear equation (3), i.e.

$$\mathbf{d}^\dagger = \mathcal{F}(\mathbf{S}^\dagger), \tag{4}$$

to find $\mathbf{S}^\dagger$. Among the numerous DI approaches to solve (4) there is a class of methods that use $\mathbf{d}^\dagger = \mathbf{u}_M^\dagger$ (suitably interpolated) instead of $\mathbf{u}$ in equation (1) to solve for $\mathbf{S}^\dagger$ (Sinkus *et al* 2005, Papazoglou *et al* 2008). Since $\mathbf{d}^\dagger$ is corrupted by noise, employing DI often requires application of filtering/smoothing on $\mathbf{d}^\dagger$ so that it can be used within the model equations that involve derivatives of the displacement field. The assumption of local homogeneity of the tissue properties Sinkus *et al* (2005) can cause significant errors in the elastograms at the boundaries between tissues having very different elasticity properties, and the errors at tissue boundaries can be

accentuated by smoothing. Further, the required differentiation can amplify imaging-related noise. Another weakness of DI is that its results are not accompanied by quantification of their uncertainty.

Iterative Inversion (II) methods for MRE (see e.g. Van Houten *et al* 1999, McGarry *et al* 2013, Ammari *et al* 2015, McGarry *et al* 2019) are the other popular approach to solving the MRE inverse problem. II methods pose the inverse problem in terms of optimisation of a least-squares functional, e.g.

$$\mathbf{S}^* = \arg\min_{\mathbf{S}\in\mathcal{S}} ||\mathbf{d}^\dagger - \mathcal{F}(\mathbf{S})||^2, \tag{5}$$

where $\mathcal{S}$ is a suitable space of admissible solutions. Because the inverse problem posed via (5) is also ill-posed as described above, an additional regularisation term, typically the squared norm of $\mathbf{S}$ on some functional space, is often included in the right hand side of (5) to ensure stability in computations of an optimiser. The review Fovargue *et al* (2018) describes various II approaches with the use of different regularisation terms and their effect on the accuracy of the reconstructions of stiffness properties.

There are three fundamental limitations in II approaches for MRE. First, since the most natural approach to compute (5) is via a gradient-based or Newton-type method, derivatives of the forward map and their corresponding adjoints must be computed. This can be a limitation when using commercial software for the forward model which may not give direct access to the linearised equations and their respective adjoints. Of course, one can always employ black-box optimisation tools based on, for example, automatic differentiation at the expense of higher computational cost and possible loss of accuracy. Second, existing II approaches for MRE do not provide a measure of uncertainty in the estimates that they produce. The third limitation is that most standard regularisation terms used in II enforce a degree of smoothness in the reconstructed stiffness properties which is not consistent with the abrupt changes that these properties display in the presence of diseased tissue (Fovargue *et al* 2018). In effect, overly-smooth reconstructed properties do not enable us to accurately estimate the interface between healthy and diseased tissue which is, in turn, detrimental to the diagnostic capability of MRE. However, there are II methods for MRE, such as the use of total variation regularisation (Zhang *et al* 2012) and modelling stiffness properties via the truncation of level-set functions (Li *et al* 2012), that take into to account possible abrupt changes in stiffness properties.

In this work we propose a new approach for solving the MRE inverse problem based on the Bayesian formulation of inverse problems and appropriate parametrisation of the considered MRE inverse problem. This approach allows us to recover sharp interfaces and to quantify uncertainties of estimates of the elastic properties of interest.

2.3. The Bayesian approach

In the Bayesian approach the stiffness properties comprised in $\mathbf{S}$ are unknown random functions on which we specify a joint *prior distribution* $\mathbb{P}(\mathbf{S})$. We assume that the observed data $\mathbf{d}^\dagger$ is a realisation of a random variable $\mathbf{d}$ defined by

$$\mathbf{d} = \mathcal{F}(\mathbf{S}) + \eta = \mathbf{u}_M + \eta, \tag{6}$$

where η is an unknown random error which follows a Gaussian distribution with zero mean and covariance Σ (that is $\eta \sim N(0, \Sigma)$) which, in practice, can be informed by the precision of the measurement device used.

The aim of the Bayesian approach is to approximate the so-called *posterior distribution* $\mathbb{P}(\mathbf{S}|\mathbf{d}^\dagger)$ which from Bayes' rule is given by Stuart (2010):

$$\mathbb{P}(\mathbf{S}|\mathbf{d}^\dagger) = \frac{1}{Z}\mathbb{P}(\mathbf{d}^\dagger|\mathbf{S})\mathbb{P}(\mathbf{S}), \tag{7}$$

where $\mathbb{P}(\mathbf{d}^\dagger|\mathbf{S})$ is the *likelihood* of the measurements $\mathbf{d}^\dagger$ and Z is the normalising constant defined by

$$Z \equiv \int_{\mathcal{S}} \mathbb{P}(\mathbf{d}^\dagger|\mathbf{S})\mathbb{P}(\mathrm{d}(\mathbf{S})). \tag{8}$$

From the Gaussian assumption on η and (6), it follows that $\mathbf{d}|\mathbf{S} \sim N(\mathcal{F}(\mathbf{S}), \Sigma)$. Hence, the likelihood is given by

$$\mathbb{P}(\mathbf{d}^\dagger|\mathbf{S}) \propto \exp\left[-\frac{1}{2}\,\|\Sigma^{-1/2}(\mathbf{d}^\dagger - \mathcal{F}(\mathbf{S}))\|^2\right]. \tag{9}$$

Once we specify the prior $\mathbb{P}(\mathbf{S})$, we can use (9) in (7) which yields the sought posterior $\mathbb{P}(\mathbf{S}|\mathbf{d}^\dagger)$ up to the normalisation constant Z. Unfortunately, since $\mathcal{F}$ is a nonlinear operator, Z cannot be computed in closed-form. Therefore, $\mathbb{P}(\mathbf{S}|\mathbf{d}^\dagger)$ must be approximated via sampling strategies, where the goal is to compute a collection of samples, $\mathbf{S}^{(j)} \sim \mathbb{P}(\mathbf{S}|\mathbf{d}^\dagger)$ $(j = 1, \ldots, J)$, from which Monte Carlo estimates can be computed. These include the sample mean, variance as well as credible intervals for the stiffness properties (at each point $\mathbf{x}$ in the computational domain Ω). Another relevant quantity is the maximum *a posteriori* (MAP) estimate which is defined by

$$\mathbf{S}_{MAP} \equiv \arg\max_{\mathbf{S}\in\mathcal{S}} \mathbb{P}(\mathbf{S}|\mathbf{d}^{\dagger}) = \arg\max_{\mathbf{S}\in\mathcal{S}} \left\{ \exp\left[-\frac{1}{2} \|\Sigma^{-1/2}(\mathbf{d}^{\dagger} - \mathcal{F}(\mathbf{S}))\|^2 \right] \mathbb{P}(\mathbf{S}) \right\}, \tag{10}$$

and which corresponds to the most likely estimate of stiffness properties of the posterior estimate.

Computing the MAP (i.e. maximising the posterior) is equivalent to an II approach (i.e. minimising the squared data misfit + regularisation term) (Kaipio and Somersalo 2005), where the regularisation term is determined by the choice of prior. The work of Mohammadi *et al* (2021) uses the so-called total variation prior to compute the MAP for a simple 2D MRE problem. Although they use the Bayesian setting to derive the approach, they do not characterise the full posterior distribution which, as stated before, is required to quantify uncertainty of all possible positions. On the other hand, a Markov Chain Monte Carlo approach was recently used in Jiang and Qian (2020) to fully characterise the posterior for inversion in a MRE setting. This approach, however, was used in a very simplified case in 2D, where the stiffness properties were piecewise constant and the interface between healthy and diseased tissue was assumed known. The Bayesian approach that we present here approximates the full Bayesian posterior of stiffness properties in 2D and 3D settings where the location of those interfaces are unknown and inferred via the proposed approach. Moreover, within each region, we allow for the inference of highly heterogeneous properties.

In the next section we parameterise the stiffness properties (the first and second Lame parameters), which is needed in order to allow for sharp discontinuities in tissues and high heterogeneity of the properties, and we specify the priors on the parameterisation introduced.

3. Parameterisation of the Bayesian inverse problem

Our main objective here is to define a parameterisation of the stiffness properties $\mathbf{S} = (G, \lambda)$. As outlined in the introduction, the second Lame parameter (or shear modulus) can be written as

$$G(\mathbf{x}) = G_s(\mathbf{x}) + iG_l(\mathbf{x}), \tag{11}$$

where $G_s(\mathbf{x})$ and $G_l(\mathbf{x})$ are the storage and loss modulus, respectively. While a similar expression holds for $\lambda(\mathbf{x})$, here we assume that the complex part of this function is negligible (see e.g Sinkus *et al* 2005 for a physical justification) and use the standard relation (Sinkus *et al* 2005)

$$\lambda(\mathbf{x}) = \frac{2\nu G_s(\mathbf{x})}{1 - 2\nu}, \tag{12}$$

in terms of the storage modulus and the Poisson's ratio denoted by ν. We employ the standard assumption of near-incompressibility of biological tissues and hence Poisson's ratio is close to that of fluid, $\nu \approx 0.5$ (Ammari *et al* 2015, McGrath 2018).

Using these assumptions, the stiffness properties (G, λ) can be characterised by the storage and loss modulus

$$\mathbf{S}(\mathbf{x}) = (G_s(\mathbf{x}), G_l(\mathbf{x})). \tag{13}$$

The framework subsequently developed is focused on inferring $\mathbf{S}$ from (13), but it is crucial to emphasise that the proposed method is quite general and can be applied to any set of physical properties of interest (e.g. ν, λ, ρ).

In the next subsection we introduce a parameterisation

$$\mathbf{S} = (G_s, G_l) = \mathcal{P}(\boldsymbol{\Theta}), \tag{14}$$

in terms of an auxiliary function $\boldsymbol{\Theta}$. This parameterisation enables one to (i) incorporate anatomical information of the tissue under investigation and (ii) infer storage and loss moduli that are highly heterogeneous and could potentially exhibit sharp discontinuities at the borders of diseased tissue. The parameterised inverse problem is then solved via the Bayesian approach in section 3.2.

3.1. Parameterisation of stiffness properties

In order to introduce the parameterisation we use the example of its application to the kidney as described in section 4, but the framework can be tailored to any other MRE settings. The kidney is modelled as being divided into two regions comprised of the medulla and cortex which we denote by Ω_m and Ω_c, respectively. The domain of interest can be represented as

$$\Omega = \Omega_b \bigcup \Omega_m \bigcup \Omega_c, \tag{15}$$

where Ω_b is the background medium surrounding the kidney (see figure 1 (left) for a 2D example).

In order to account for (and subsequently infer) the presence of diseased tissue (e.g. a cystic lesion within the kidney), we introduce the set $\Omega_d \subset \Omega_c \cup \Omega_m$. Let us now express the storage and loss moduli in terms of a collection of functions defined on each of the sections of the domain of interest. More specifically, we assume that for the multi-index $\alpha \in \{s, l\}$:

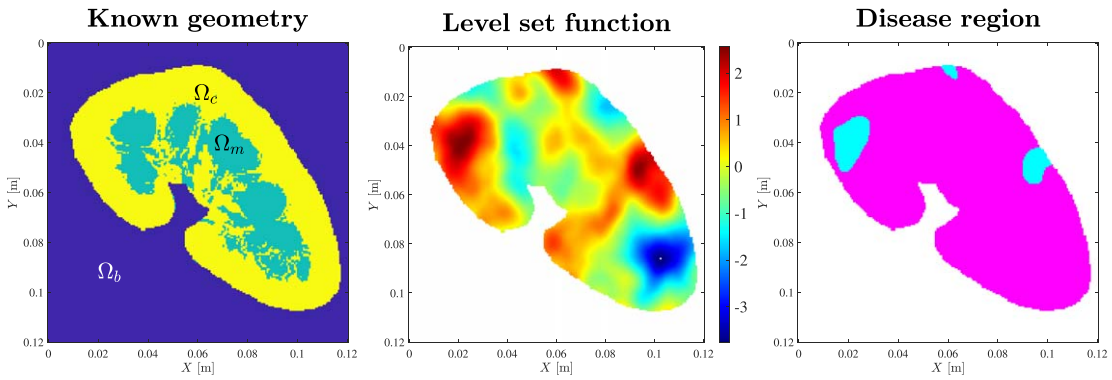

Figure 1. 2D Setting. Left: setting for the known geometry. Middle: example of a random level-set function (dimensionless). Right: unknown diseased (blue) region characterised via the truncation of the level-set function.

$$G_\alpha(\mathbf{x}) = \begin{cases} G_{\alpha,c}(\mathbf{x}), & \mathbf{x} \in \Omega_c \setminus \Omega_t, \\ G_{\alpha,m}(\mathbf{x}), & \mathbf{x} \in \Omega_m \setminus \Omega_t, \\ G_{\alpha,d}(\mathbf{x}), & \mathbf{x} \in \Omega_d, \\ G_{\alpha,b}(\mathbf{x}), & \mathbf{x} \in \Omega_b. \end{cases} \tag{16}$$

For simplicity we assume that we know the geometry of the cortex and the medulla regions, Ω_c and Ω_m, while we focus on inferring the diseased region Ω_d. Although the extent of the diseased region Ω_d may be also available from the MRI measures (e.g. T1-weighted or T2-weighted MRI images, or quanttative maps), we keep Ω_d as unknown and hence demonstrate the flexibility of the proposed EKI approach to infer the geometry of unknown regions. To this end, we assume that Ω_d is unknown and parameterised via the level-set approach (Iglesias *et al* 2016, Dunlop *et al* 2017). That is

$$\Omega_d = \{x \in \Omega_c \bigcup \Omega_m \colon f(\mathbf{x}) \leqslant \kappa\}, \tag{17}$$

where f is the level-set function and κ is a threshold. We note that the κ-level corresponds to the interface between diseased and healthy kidney tissue. While κ can be subsequently inferred as part of the unknown, here we simply assume it is fixed and specified beforehand.

For the experiments of subsequent sections, the geometry of the regions Ω_c and Ω_m is assumed known from, for example, segmented images obtained via thresholding the signal of a T1-weighted MR image, here an *ex vivo* kidney specimen. The practical case in which segmentation errors are present (and so uncertainties in the geometry Ω_c and Ω_m) can also be addressed within the proposed EKI level-set framework. This is done in section 4 of the *supplementary material*.

It is worth mentioning that the use of level sets to model discontinuous stiffness properties in MRE has been also proposed in (Li *et al* 2012) although in a framework that (i) can only be applied to piece-wise constant elastic properties and (ii) is purely deterministic and hence does not provide uncertainty measures of the produced estimates.

Expression (16) with Ω_d characterised via (17) defines a map, $\mathcal{Q}$, such that

$$(\{G_{s,\beta}\}_{\beta\in\mathcal{B}}, \{G_{l,\beta}\}_{\beta\in\mathcal{B}}, f) \longrightarrow \mathbf{S} = (G_s, G_l) = \mathcal{Q}(\{G_{s,\beta}\}_{\beta\in\mathcal{B}}, \{G_{l,\beta}\}_{\beta\in\mathcal{B}}, f), \tag{18}$$

where we have introduced the set of indices $\mathcal{B} \equiv \{b, c, d, m\}$. The map $\mathcal{Q}$ enables us to parameterise the unknown stiffness properties $\mathbf{S} = (G_s, G_l)$ in terms of the loss and storage modulus on each sub-region of interest (i.e. $G_{\alpha,\beta}$), as well as the level-set function, f, which determines possible presence of an unknown disease region. We further characterise each of these functions via Random Fields (RFs) that we generate using a Whittle–Matern parameterisation, $\mathcal{P}^{WM}$, which we describe in appendix A. More specifically, we assume that (for $\beta \in \mathcal{B}$)

$$G_{s,\beta} = \exp[\mathcal{P}^{WM}(\boldsymbol{\Theta}_{s,\beta})], \quad G_{l,\beta} = \exp[\mathcal{P}^{WM}(\boldsymbol{\Theta}_{l,\beta})], \quad f = \mathcal{P}^{WM}(\boldsymbol{\Theta}_f), \tag{19}$$

where $\boldsymbol{\Theta}_{l,\beta}$, $\boldsymbol{\Theta}_{s,\beta}$ and $\boldsymbol{\Theta}_f$ are additional auxiliary functions and parameters that we comprise in

$$\boldsymbol{\Theta}(\mathbf{x}) = (\{\boldsymbol{\Theta}_{s,\beta}(\mathbf{x})\}_{\beta\in\mathcal{B}}, \{\boldsymbol{\Theta}_{l,\beta}(\mathbf{x})\}_{\beta\in\mathcal{B}}, \boldsymbol{\Theta}_f(\mathbf{x})). \tag{20}$$

Exponential functions in (19) are used to ensure the values of loss/storage obtained for any choice of $\boldsymbol{\Theta}$ are always positive. Let us note that if we denote by $\mathcal{P}$ the composition of (18) with (19), we can succinctly write the parameterisation of $\mathbf{S}$ as in (14).

The formulation of $\mathcal{P}^{WM}$ as well as the expression for $\boldsymbol{\Theta}$ are discussed in appendix A. It is worth mentioning that when $\boldsymbol{\Theta}$ is selected as discussed in appendix B, then the $\log(G_{s,\beta})$'s, $\log(G_{l,\beta})$'s and f that we obtain via (11)

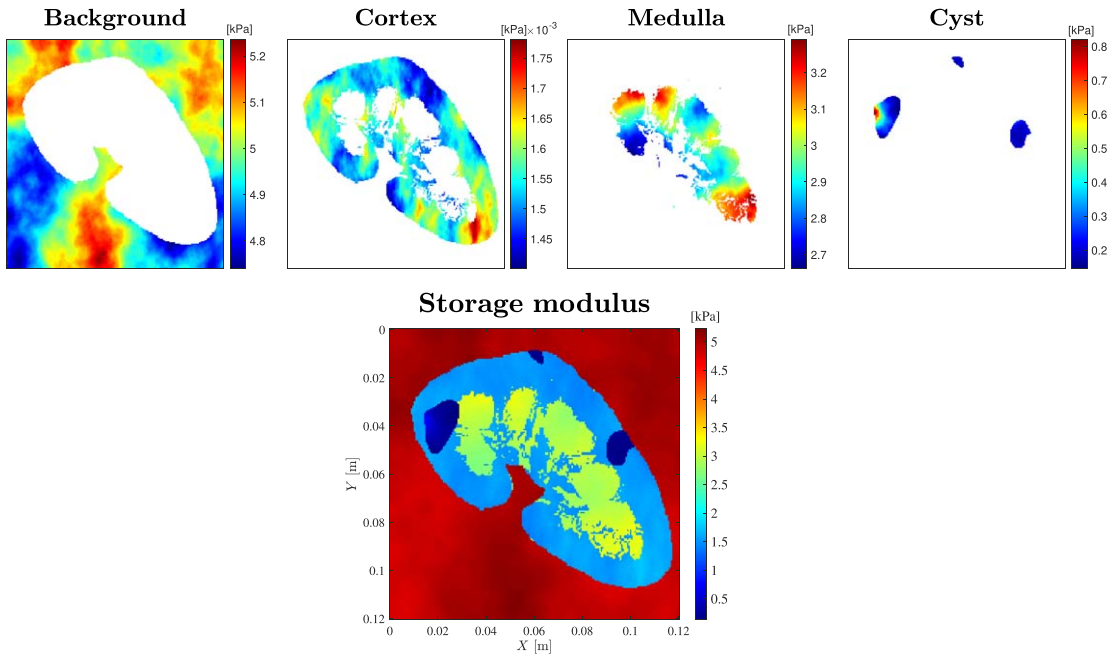

Figure 2. Parameterisation. Top: examples of random fields used to characterise the storage modules [kPa] on each anatomical region. Bottom: a random storage modulus [kPa] computed with the proposed parameterisation.

are Gaussian RFs which have been widely used to model spatial variability of physical properties in many other applications (Matérn 1986, Stein 1999, Lasanen *et al* 2014).

Before we proceed to discuss the Bayesian approach that we employ to infer $\boldsymbol{\Theta}$ (and hence G_s and G_l via (14)), let us illustrate the proposed parameterisation using the 2D geometry from figure 1 (left). We reiterate that the geometry of the cortex and medulla are assumed to be available from prior information, but the approach can be extended to infer these regions within the EKI algorithm. In the middle panel of figure 1 we provide an example of a Gaussian RF that we generate as described in appendix B and that we use as the level-set function, f, in definition (17). In addition, we employ a threshold value $\kappa = 1.75$ to obtain the cystic region, Ω_d, that we display in blue colour in figure 1 (right). Additional Gaussian RFs are generated to characterise the (log) storage modulus within each of the regions. These fields, with plots shown in figure 2, correspond to (from left to right) $G_{s,b}$, $G_{s,c}$, $G_{s,m}$ and $G_{s,d}$. If we use these functions, together with the level-set function from figure 1 in expression (16), we finally obtain the storage modulus from figure 2 (bottom).

3.2. The re-parameterised inverse problem

In appendix B we discuss in detail the prior that we choose for the parameters comprised in $\boldsymbol{\Theta}$. The prior on the stiffness properties is defined by

$$\mathbb{P}(\mathbf{S}) \equiv \mathcal{P}^{\#}\mathbb{P}(\boldsymbol{\Theta}),$$

where $\mathcal{P}^{\#}\mathbb{P}$ denotes the push-forward measure of $\mathbb{P}$ under the parameterisation map $\mathcal{P}$. In practical terms this means that the prior, $\mathbb{P}(\mathbf{S})$, that we define on $\mathbf{S}$ (see (13)) can be simply characterised using samples from $\mathbb{P}(\boldsymbol{\Theta})$ mapped under $\mathcal{P}$, i.e.

$$\boldsymbol{\Theta}^{(j)} \sim \mathbb{P}(\boldsymbol{\Theta}) \Longrightarrow \mathbf{S}^{(j)} = (G_s^{(j)}, G_l^{(j)}) = \mathcal{P}(\boldsymbol{\Theta}^{(j)}) \sim \mathbb{P}(\mathbf{S}).$$

We can apply the Bayesian framework on the problem parameterised in terms of $\boldsymbol{\Theta}$. In other words, we seek the posterior

$$\mathbb{P}(\boldsymbol{\Theta}|\mathbf{d}^{\dagger}) = \frac{1}{Z}\exp\left[-\frac{1}{2}\,\|\Gamma^{1/2}(\mathbf{d}^{\dagger} - \mathcal{H}(\boldsymbol{\Theta}))\|^2\right]\mathbb{P}(\boldsymbol{\Theta}), \tag{21}$$

where Z is the normalisation constant similar to the one defined in (8) and where we have defined the parameter-to-measurements map

$$\mathcal{H} \equiv \mathcal{F}\circ\mathcal{P}. \tag{22}$$

In (21) we have used the likelihood implied by (6) with $\mathbf{S}$ parameterised as in (14). It can be shown (e.g. see Dunlop *et al* 2020) that

$$\mathbb{P}(\mathbf{S}|\mathbf{d}^\dagger) = \mathcal{P}^{\#}\mathbb{P}(\boldsymbol{\Theta}|\mathbf{d}^\dagger).$$

Hence, if we compute a sampling approximation of the Bayesian posterior on the re-parameterised problem, samples from the posterior on the stiffness properties can be computed by means of mapping those samples under $\mathcal{P}$.

In this work we approximate the posterior $\mathbb{P}(\boldsymbol{\Theta}|\mathbf{d}^\dagger)$ by means of the EKI algorithm initially introduced in Iglesias *et al* (2013), Iglesias (2016) for solving generic inverse problems. Here we briefly introduce the algorithm and refer the reader to the recent work of Iglesias and Yang (2021) for further details and motivation of EKI from the Bayesian tempering setting for inverse problems.

The underlying framework of the EKI algorithm that we display in algorithm 1 is to approximate, via an ensemble of J particles/samples, a sequence of intermediate distributions between prior and posterior:

$$\mathbb{P}_0(\boldsymbol{\Theta}) \equiv \mathbb{P}(\boldsymbol{\Theta}) \;\rightarrow\; \mathbb{P}_1(\boldsymbol{\Theta}) \rightarrow \cdots \rightarrow \mathbb{P}_{N+1}(\boldsymbol{\Theta}) \equiv \mathbb{P}(\boldsymbol{\Theta}|\mathbf{d}^\dagger).$$

Each of these intermediate distributions $\mathbb{P}_n(\boldsymbol{\Theta})$ is approximated with a Gaussian distribution which is, in turn, characterised by an ensemble of particles $\{\boldsymbol{\Theta}^{(j)}\}_{j=1}^J$.

The EKI algorithm starts with an ensemble of samples from the prior which are updated at each iteration of the algorithm. The number of iterations is determined adaptively as suggested in Iglesias and Yang (2021) in order to provide stable transitions between intermediate distributions. Upon convergence, the EKI produces an ensemble that approximates the posterior.

Algorithm 1. EKI algorithm

Input : (1) $\mathbf{d}^\dagger$: measurements; (2) Γ measurements' error covariance;
(3) $\{\boldsymbol{\Theta}_0^{(j)}\}_{j=1}^J$: initial ensemble consisting of samples from the prior.
Output: $\{\boldsymbol{\Theta}^{(j)}\}_{j=1}^J$: posterior ensemble.
Set $s_0 = 0$
while $s_n < 1$ **do**
(1) **Prediction step.** Evaluate

$$\mathcal{H}_n^{(j)} = \mathcal{H}(\boldsymbol{\Theta}_n^{(j)}) = \mathcal{F}(\mathcal{P}(\boldsymbol{\Theta}_n^{(j)})), \quad j \in \{1, \ldots, J\}, \tag{23}$$

and define $\overline{\mathcal{H}}_n = \frac{1}{J}\sum_{j=1}^J \mathcal{H}_n^{(j)}$.
(2) **Compute regularisation parameter** α_n**:**

$$\alpha_n^* = \frac{1}{M}\frac{1}{J}\sum_{j=1}^J \|\Gamma^{-1/2}(\mathbf{d}^\dagger - \mathcal{H}_n^{(j)})\|^2$$

if $s_n + \frac{1}{\alpha_n^*} \geqslant 1$ **then**
Set $\alpha_n = \frac{1}{1-s_n}$, $s_{n+1} = 1$;
else
Set $\alpha_n = \alpha_n^*$, $s_{n+1} = s_n + \frac{1}{\alpha_n}$.
end 3
Analysis step. Define $C_n^{\mathcal{H}\mathcal{H}}$, $C_n^{\Theta\mathcal{H}}$ by

$$C_n^{\mathcal{H}\mathcal{H}} = \frac{1}{J-1}\sum_{j=1}^J (\mathcal{H}_n^{(j)} - \overline{\mathcal{H}}_n)(\mathcal{H}_n^{(j)} - \overline{\mathcal{H}}_n)^T,$$
$$C_n^{\Theta\mathcal{H}} = \frac{1}{J-1}\sum_{j=1}^J (\boldsymbol{\Theta}_n^{(j)} - \overline{\boldsymbol{\Theta}}_n)(\mathcal{H}_n^{(j)} - \overline{\mathcal{H}}_n)^T.$$

where $\overline{\boldsymbol{\Theta}}_n = \frac{1}{J}\sum_{j=1}^J \boldsymbol{\Theta}_n^{(j)}$, and update each ensemble member:

$$\boldsymbol{\Theta}_{n+1}^{(j)} = \boldsymbol{\Theta}_n^{(j)} + C_n^{\Theta\mathcal{H}}(C_n^{\mathcal{H}\mathcal{H}} + \alpha_n\Sigma)^{-1}(\mathbf{d}^\dagger - \mathcal{H}_n^{(j)} + \sqrt{\alpha_n}\eta_n^{(j)}),$$
$$j \in \{1, \ldots, J\}, \tag{24}$$

where $\eta_n^{(j)} \sim N(0, \Sigma)$.
$n + 1 \rightarrow n$
end

At every iteration of the scheme, the main computational cost of EKI is in the prediction step which consists in evaluating the parameter-to-measurement map $\mathcal{H}$ for each of the ensemble members $\boldsymbol{\Theta}^{(j)}$. This evaluation involves (see equation (23)) (i) computing the stiffness properties $\mathbf{S}^{(j)} = \mathcal{P}(\boldsymbol{\Theta}^{(j)})$ and (ii) tissue displacements $\mathcal{H}^{(j)} = \mathcal{F}(\mathbf{S}^{(j)})$. The latter involves solving equation (1) while the former, as discussed in appendix A, requires solving one additional PDE for each of the storage and the loss modulus defined on the constituents of Ω as well as the level-set function. If we denote by $\mathcal{C}_H$ the total cost incurred by one evaluation of the parameter-to-measurement map $\mathcal{H}$ and N_{it} the iterations needed to achieve convergence, then the total cost of EKI is

$$\mathcal{C}_{EKI} = N_{it} J \mathcal{C}_H.$$

To the best of our knowledge, this work constitutes the first application of EKI to time-harmonic MRE. For quasi-static MRE, however, the recent work of Napoli *et al* (2021) used a Kalman filter algorithm to assimilate data from robotic sensors. Their algorithmic framework was entirely different to the Bayesian inversion setting proposed here and without the use of the level-set based parameterisation of stiffness properties.

3.3. Measures of uncertainty

As we discussed earlier, samples from $\boldsymbol{\Theta}$ (either from the prior or the approximate posterior computed via EKI) can be mapped, via (14), into samples of the corresponding distributions of the physical properties G_s and G_l. In other words, we have the following samples

$$(G_s^{(j)}, G_l^{(j)}) = \mathcal{P}(\boldsymbol{\Theta}^{(j)}), \qquad j = 1, \ldots, J, \tag{25}$$

from which we may compute the sample mean and variance (for $\alpha \in \{s, l\}$) defined by

$$\overline{G}_\alpha(\mathbf{x}) = \frac{1}{J}\sum_{j=1}^{J} G_\alpha^{(j)}(\mathbf{x}) \quad \text{and} \quad \sigma^2_{G_\alpha}(\mathbf{x}) = \frac{1}{J-1}\sum_{j=1}^{J}(G_\alpha^{(j)}(\mathbf{x}) - \overline{G}_\alpha(\mathbf{x}))^2, \tag{26}$$

respectively. When we use samples of the prior (or posterior), the expressions in (26) converge, as $J \to \infty$, to the mean and the variance of the prior (or posterior) distribution. We reiterate that EKI produces a Gaussian approximation of the posterior. Although the posterior is, in general, non-Gaussian, there is numerical evidence from other applications that EKI can accurately approximate the posterior mean and variance (Iglesias 2014, Iglesias *et al* 2018, Matveev *et al* 2021).

Similarly, for each $\mathbf{x} \in \Omega$ we use our samples, $\{G_\alpha^{(j)}(\mathbf{x})\}_{j=1}^{J}$ to approximate equal tail $(1-\alpha)100\%$ credible intervals (Gelman *et al* 2004), i.e. intervals that contain the unobserved properties (at that given $\mathbf{x}$) with $(1-\alpha)$ 100% posterior probability. In addition, the level-set parameterisation of Ω_d (see equation (17)) allows us to compute the posterior probability that a point $\mathbf{x}$, within the tissue, belongs to the diseased tissue region. More specifically, we define

$$P(\mathbf{x}) \equiv \mathbb{P}(f(\mathbf{x}) > \kappa), \tag{27}$$

which can be approximated by

$$P(\mathbf{x}) \simeq \frac{1}{J}\sum_{j=1}^{J} \mathbb{1}_{\Omega^{(j)}}(\mathbf{x}), \tag{28}$$

where $\mathbb{1}_{\Omega^{(j)}}(\mathbf{x})$ is the indicator function of $\Omega^{(j)}$ defined by

$$\mathbb{1}_{\Omega^{(j)}}(\mathbf{x}) = \begin{cases} 1 & x \in \ \Omega^{(j)}, \\ 0 & x \notin \ \Omega^{(j)}, \end{cases} \tag{29}$$

and $\Omega^{(j)}$ is given by (17) with $f^{(j)} = \mathcal{P}^{WM}(\boldsymbol{\Theta}_f^{(j)})$ instead of f. Here $\boldsymbol{\Theta}_f^{(j)}$ is one of the components of the jth posterior ensemble of the parameter $\boldsymbol{\Theta}^{(j)}$ (see equation (20)). Further details on how to arrive at (28) from (27) can be found in Matveev *et al* (2021), where a similar quantity was defined for the probability defects in composites materials.

4. Numerical experiments

In this section we test the proposed Bayesian inversion algorithm using synthetic data. The use of synthetic data allows us to demonstrate accuracy and the cost of the method in a transparent way.

In section 4.1 we specify settings of the forward model for the experiments. Results of 2D experiments are presented in section 4.2 while in section 4.3 results of 3D tests are given.

For the 2D case, in the *supplementary material* we look in depth on how the ensemble size (section 1 of the *supplementary material*), the level of measurement errors (section 2 of the *supplementary material*), and the choice of prior (section 3 of the *supplementary material*) affect the outcomes of our inversion algorithm. Furthermore, in section 4 of *supplementary material* the proposed EKI-based inversion for MRE is extended to account for segmentation errors/uncertainties, which is accompanied by illustrative numerical experiments. In section 5 of *supplementary material* we compare the performance of the proposed level-set based EKI inversion with the standard implementation of the LFE method encoded in the software package MRE Wave provided by the Mayo Clinic, and we show the superiority of the proposed method.

The code for the 2D experiments can be downloaded from https://github.com/Marco-Iglesias-Nottingham/EKI-for-MRE.

4.1. Forward modelling settings

The forward model that we use for our numerical experiments is given by equation (1), where G and λ are parameterised in terms of G_s and G_l via (11)–(12). For all our experiments MRE vibration frequency was 60 Hz, and parameters employed were:

$$\omega = 2\pi \ 60 \text{ rad}, \quad \rho = 10^3 \text{ kg m}^{-3} \quad \text{and} \quad \nu = 0.499.$$

We consider the following boundary conditions (see e.g. Ciarlet 1988, Alberty *et al* 2002):

$$[\lambda(\nabla \cdot \mathbf{u})\mathbf{I} + G[\nabla \mathbf{u} + (\nabla \mathbf{u})^T]]\mathbf{n} = 0, \qquad \text{on } \partial\Omega_N \tag{30}$$

$$\mathbf{M}\mathbf{u} = \mathbf{u}_D, \qquad \text{on } \partial\Omega_D, \tag{31}$$

where it is assumed that the boundary of Ω, denoted by $\partial\Omega$, can be written as $\partial\Omega = \partial\Omega_N \cup \partial\Omega_D$, on which we may impose natural (for simplicity traction-free) and Dirichlet boundary conditions, respectively. In equation (31) $\mathbf{M}$ is a $d \times d$ matrix that we use to constrain some of the components of $\mathbf{u}(\mathbf{x})$ on $\partial\Omega_D$, while the other components satisfy the traction-free condition.

It is worth noticing that these boundary conditions are chosen for the sake of simplicity in demonstrating the proposed approach for Bayesian inversion in MRE. We note that in real settings, boundary conditions will be unknown but observed displacements (from MRI measurements) could be used to specify Dirichlet boundary values on a specified boundary. In addition, while we have selected a common choice of 60 Hz for the value of frequency, since the proposed EKI approach uses the (forward) viscoelastic model in a black box fashion, different choices of frequency should not, in principle, affect the outcomes of the inversion provided (i) the same level of noise is employed and (ii) tissue displacements are observed on a sufficiently dense grid (see equation (2)) so the the underlying wavelength can be resolved.

For the 2D experiments of section 4.2 the domain is $\Omega = [0,\ 0.12 \text{ m}]^2$ with Dirichlet boundary, $\partial\Omega_D$, given by the horizontal edge at $y = 0.12$ m. For the 3D experiments reported in section 4.3 we use $\Omega = [-4.57 \times 10^{-2} \text{ m}, 4.57 \times 10^{-2} \text{ m}] \times [-6.9 \times 10^{-2} \text{ m}, 6.9 \times 10^{-2} \text{ m}] \times [0, 3.0 \times 10^{-2} \text{ m}]$ and $\partial\Omega_D$ given by two faces of the hexahedron; one face defined by $y = 6.9 \times 10^{-2}$ m and the other one by $x = -4.57 \times 10^{-2}$ m. For both 2D and 3D settings, we (i) define the matrix $\mathbf{M}$ so that only the y-component is fixed on $\partial\Omega_D$ and (ii) set the magnitude of $\mathbf{u}_D$ equal to 10^{-6} m.

The forward modeling setting is implemented using Finite Elements Method via the MATLAB toolbox `pdetool`. We use the built-in command `generateMesh` to produce a mesh with 40 692 and 47 646 elements in 2D and 3D, respectively. The 2D mesh that we use is too fine to be visualised but the 3D mesh that we use is displayed in the *supplementary material*. As discussed in appendix A.1, stiffness properties (loss and storage modulus) are produced on regular grids with 100×100 and $70 \times 70 \times 30$ cells for the 2D and 3D case, respectively. Hence, the resolution of the images of stiffness properties that we produce in 2D and 3D are 1.2 mm $\times$ 1.2 mm and 1.97 mm $\times$ 1.36 mm $\times$ 1 mm, respectively. These properties are then interpolated on the mesh computed by `generateMesh` to compute displacements via `pdetool`. It is important to emphasize that the resolution used for the stiffness properties is not the same as the resolution of the (synthethic) MRE displacement data that we generate for our experiments (see below).

4.2. 2D Experiment

In figure 3 we plot the 'ground truth' fields for the storage, $G_s^\dagger$, and loss modulus, $G_l^\dagger$, for our experiments. The diseased region consists of two randomly distorted disks within the cortex. The (log) storage/loss modulus within each region are Gaussian RFs that we produce with the procedure outlined in section 3.1. The specific choice of parameters for the Whittle–Matern parameterisation (see appendix A) are shown in table 1. The mean of the random field for each region (i.e. m in table 1) were chosen to approximate the biomechanical properties of the kidney tissue types (medulla and cortex) and cysts along with the background tissue surrounding the kidney in the human body. Additional parameters (e.g. length scales variance and smoothness) have been chosen in order to add a realistic degree of the spatial variability to stiffness properties of anatomical regions of the kidney.

To generate synthetic data, we use $G_s^\dagger$ and $G_l^\dagger$ to solve the forward model with the settings described in section 4.1. The displacement fields, denoted by $\mathbf{u}^\dagger$ are evaluated on a regular grid with $M = 80 \times 80 = 6400$ equally distributed points. Therefore, the resolution of the synthetic MRE data for the 2D case is 1.5 mm $\times$ 1.5 mm. Following the notation from section 2.1, $\mathbf{u}_M^\dagger = (u_1^\dagger, \ldots, u_M^\dagger)$ corresponds to the noise-free predictions of the forward model for the ground truth. We note that $\mathbf{u}_M^\dagger$ comprises evaluation of the x (horizontal) and y (vertical) components of real and imaginary parts of $\mathbf{u}_M^\dagger$. Synthetic data $\mathbf{d}^\dagger$ are obtained by adding $\eta^\dagger \sim N(0, \Sigma)$ to $\mathbf{u}_M^\dagger$ (see equation (3)). For simplicity we consider the case in which Σ is a diagonal matrix $\Sigma = \text{diag}(\Sigma_1, \ldots, \Sigma_M)$ which corresponds to the situation where observational errors are uncorrelated. Furthermore, we model the variance of each measurement via

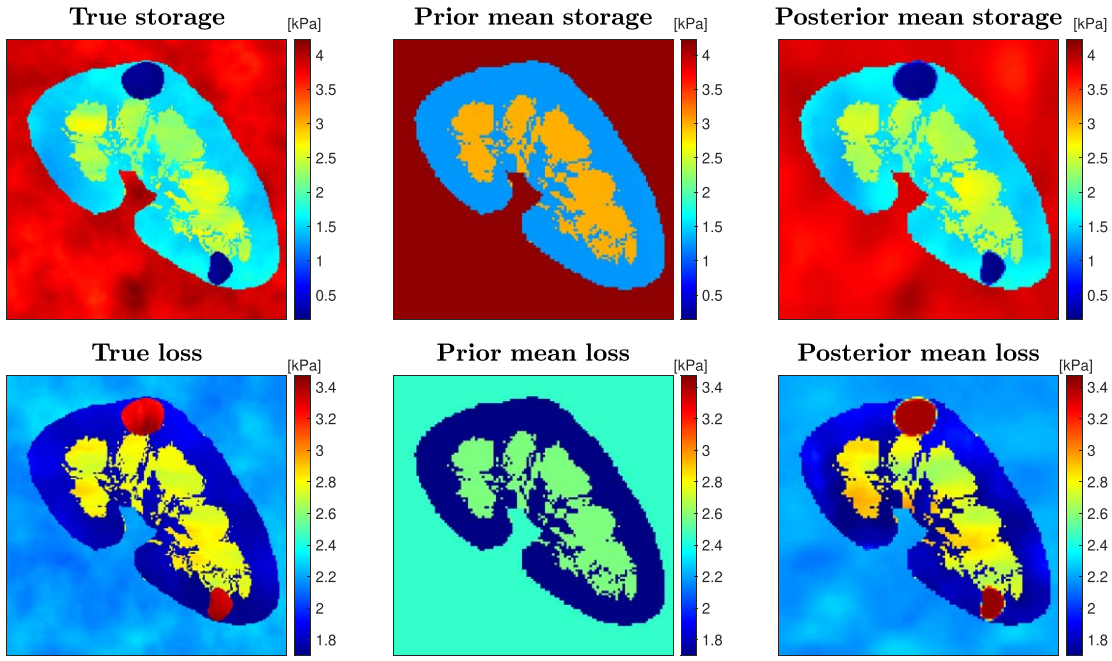

Figure 3. Comparing the ground truth (left) with prior (centre) and posterior (right) means for storage (top), $G_s^\dagger$ [kPa] and loss (bottom) modulus, $G_l^\dagger$ [kPa].

Table 1. Parameters for the Whittle–Matern parameterisation of each $\log G^\dagger_{\alpha,\beta}$ (loss: $\alpha = l$, storage: $\alpha = s$) defined on each region, i.e. background ($\beta = b$), cortex ($\beta = c$), medulla ($\beta = m$) and disease ($\beta = d$).

α	β	σ [log (Pa)]	m [kPa]	L_x [m]	L_y [m]	ζ
s	b	0.04	3.8	1.2×10^{-2}	1.2×10^{-2}	2
s	c	0.1	1.6	1.8×10^{-2}	1.8×10^{-2}	2
s	m	0.05	2.5	4.2×10^{-3}	1.2×10^{-2}	2
s	d	0.2	0.212	1.2×10^{-2}	6.0×10^{-3}	2
l	b	0.02	2.2	1.2×10^{-2}	1.2×10^{-2}	2
l	c	0.03	1.8	1.8×10^{-2}	1.8×10^{-2}	2
l	m	0.015	2.8	4.2×10^{-3}	1.2×10^{-2}	2
l	d	0.035	3.25	1.2×10^{-2}	6.0×10^{-3}	2

$$\Sigma_m = (\alpha_1 |u_m^\dagger|)^2 + (\alpha_2 |\max\{u_m^\dagger\}_{m=1}^M - \min\{u_m^\dagger\}_{m=1}^M|)^2 \qquad m = 1, \ldots, M, \tag{32}$$

where α_1 and α_2 are parameters that we select to control the level of noise/measurement error that we add to our synthetic data. The first part of the error in the right-hand side of (32) enables us to specify noise as a percentage of the noise-free measurements. With the second term, we control an additive error proportional to the amplitude of the deformation.

In practical settings the parameters α_1 and α_2 can be informed by the precision of the measurement device (i.e. the MRI scanner) used. When these parameters are unknown, they can be included as part of the unknown that we wish to infer within the Bayesian framework (Gelman *et al* 2004). For simplicity, here we assume that these parameters are known. For the first set of experiments we choose $\alpha_1 = 0.025$ which corresponds to 2.5% of noise (or SNR = 40) and a very small $\alpha_2 = 5 \times 10^{-4}$ simply to avoid zero entries in the diagonal of Σ. With these choices of measurement error, the horizontal and vertical components of the real part (denoted by $\mathbf{d}^{\mathrm{x}}_{\mathrm{re}}$ and $\mathbf{d}^{\mathrm{y}}_{\mathrm{re}}$) as well as the imaginary part ($\mathbf{d}^{\mathrm{x}}_{\mathrm{im}}$ and $\mathbf{d}^{\mathrm{y}}_{\mathrm{im}}$) of the synthetic data $\mathbf{d}^\dagger$ are displayed in figure 4.

The selection of the prior for the parameter $\boldsymbol{\Theta}$ that we infer via the EKI algorithm is discussed in appendix B, where we also look at how to select this prior in practical settings. We use (24) to map prior samples from $\boldsymbol{\Theta}$ into samples for the prior on the storage and loss modulus (G_s, G_l) and show the plot of a few representative samples in the top and top-middle rows of figure 5. A few prior samples for the loss modulus are presented in the *supplementary material*. For the sake of validation, we have selected the prior on $\boldsymbol{\Theta}$ in such a way that the corresponding prior samples for (G_s, G_l) display a very large degree of variability and, hence, uncertainty in the

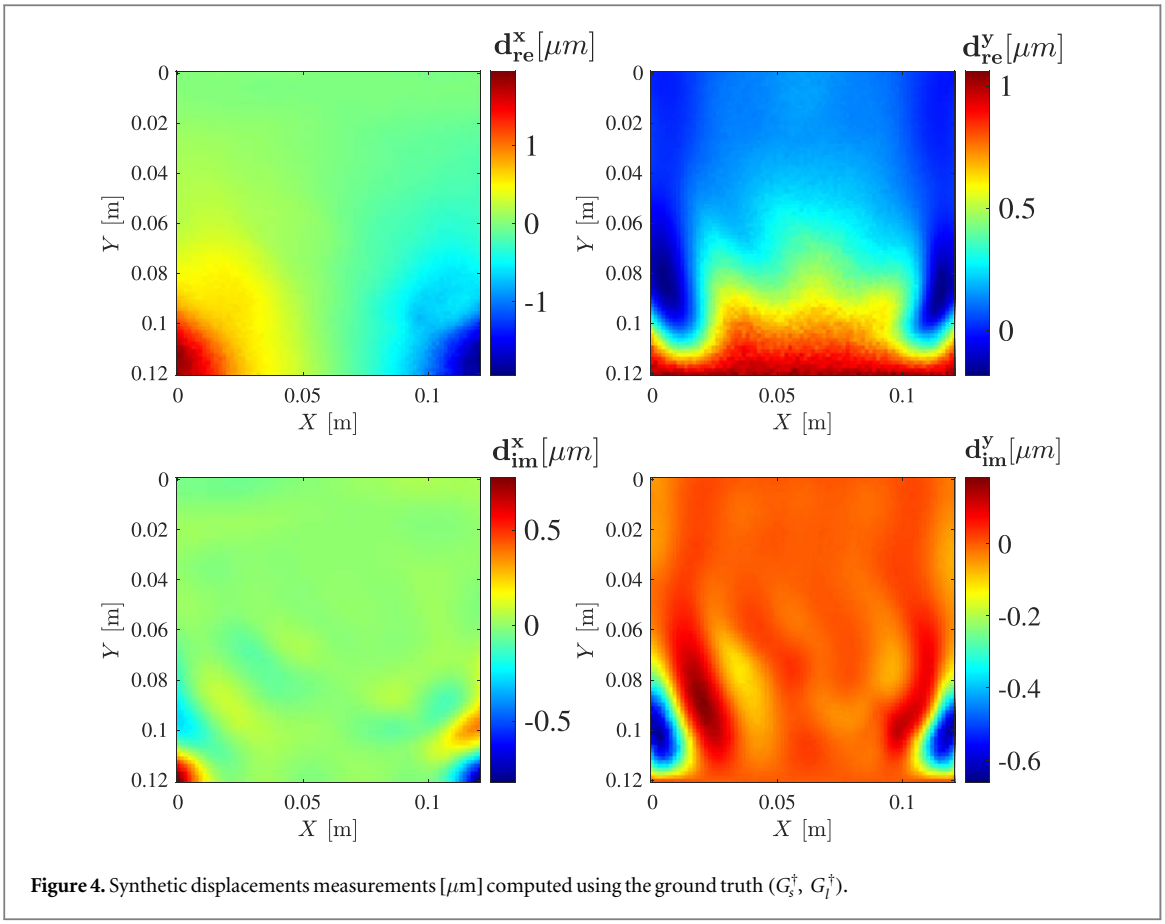

Figure 4. Synthetic displacements measurements [μm] computed using the ground truth ($G_s^\dagger$, $G_l^\dagger$).

values of the storage/loss modulus of each region. Importantly, our aim is to include as little prior information as possible of the true storage/loss in order to validate the capability of EKI to capture the true values within the posterior uncertainty measures. Similarly, our prior accounts for a substantial range of scenarios for the presence of disease. Indeed, figure 5 shows that the prior include samples in which the disease region has various random shapes and locations as well as samples without disease.

Using an ensemble size of $J = 10^4$ samples from the prior, we compute the mean and the variance for the loss/storage modules via (26). The prior means for the loss/storage modulus are shown in the middle panels of figure 3. The log of the prior variance for loss and storage are shown in the top-left and top-middle of figure 6. Using (28) we compute the (prior) probability of a presence of disease at any given point (x, y). The corresponding probability map is shown in the top-right panel of figure 6 where we see that our prior assumption is that, in average (from all samples), no disease is present.

We apply the EKI framework from algorithm 1 using synthetic data, the prior ensemble (with $J = 10^4$ particles), and covariance Σ as described earlier. Convergence is achieved after 7 iterations. Some of the posterior samples for the storage are shown in the middle-bottom and bottom rows of figure 5. Some posterior samples for the loss modulus are displayed in the *supplementary material*. We can visually appreciate that these posterior samples have all similar mean storage/loss modulus to that of the true storage/loss (see figure 3). There is, however, some noticeable differences mainly in the disease region displayed on each of these posterior samples. While all capture well the true disease region (see figure 3), some display a disease region in the upper-left area of the cortex which does not appear in the true storage/loss. Nonetheless, the posterior mean and variance, displayed in the right panels of figure 3, reveal that algorithm 1 recovered the kidney properties very accurately. The posterior means for the storage and loss provide very good estimates of the true storage and loss, respectively (see left panels of figure 3). Moreover, for most regions within the domain, the posterior variances for storage/loss are substantially smaller than those of the prior. The higher posterior variance around tissue boundaries means that there is a larger uncertainty in recovery of the properties near the boundaries and of the boundaries themselves. Also, the higher variance in the upper-left section of the cortex is consistent with the small disease region that some of the posterior samples show within this area of the kidney.

The posterior diseased tissue probability, given in the bottom-right panel of figure 6, shows remarkably accurate prediction of the true disease region within the high posterior probability area predicted by the EKI algorithm. Note that the disease region in the upper-left of the cortex that is shown in some of the samples, but that is not present in the true storage/loss, has only very low posterior probability values.

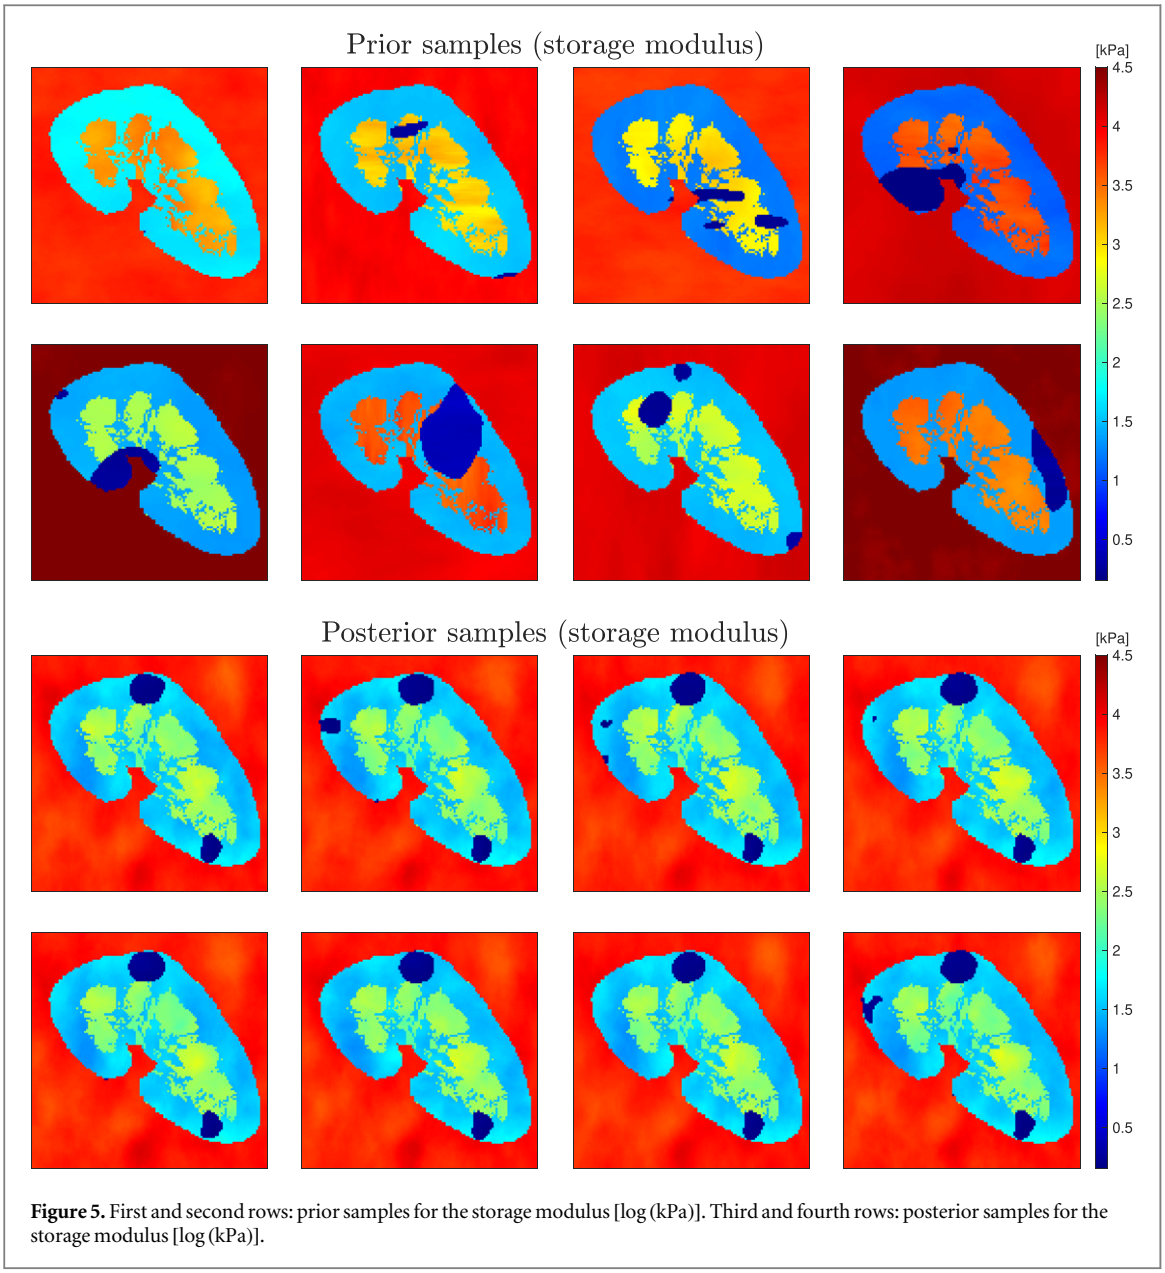

Figure 5. First and second rows: prior samples for the storage modulus [log (kPa)]. Third and fourth rows: posterior samples for the storage modulus [log (kPa)].

Figure 7 shows the graph of the true storage modulus computed for $(x, y) \in \Omega$ with fixed $y = 0.065$ m. We superimpose the graph of the mean as well as the %95 credible interval (CI) computed using the prior samples. We can see that while most of the true values are contained within the prior CI, there is substantial amount of prior uncertainty in the storage modulus. In contrast, the corresponding posterior CI computed using the EKI samples (right panel of figure 7) is concentrated around the true values. Moreover, the posterior mean provides a very good approximation to the truth as the middle-top panel of figure 3 already confirmed. In addition, in the right panel of figure 7 we have highlighted (in yellow) the area corresponding to the presence of the disease region. We recall that this region appears within the cortex. Although there is no knowledge of the presence of the disease region (cyst) within the prior values of storage modulus assigned to the cortex, we can observe that the level-set parameterisation is able to reconstruct, within the posterior CI, the sharp interface between the values of healthy cortex and those for the disease region. This showcases the advantages of the proposed EKI-based approach with respect to the existing approaches that often produce overly-smoothed stiffness properties with no well-defined boundaries between healthy and diseased tissue.

While the main focus of this paper is to infer spatial variability of the storage and loss modulus (i.e. $G_{\alpha,\beta}$) within each sub-region of Ω, it is useful to estimate the average/mean values of each of these properties. As discussed in appendix A, the average values of the $G_{\alpha,\beta}$'s are among the set of parameters comprised in $\boldsymbol{\Theta}$ which we infer via the EKI algorithm. We have selected uniform priors on these values (see appendix B) on wide intervals that contain the corresponding average values for the storage/loss modulus of the ground truth (see values of m in table 1). The posterior distribution for each of these values is one of the direct outcomes of the EKI

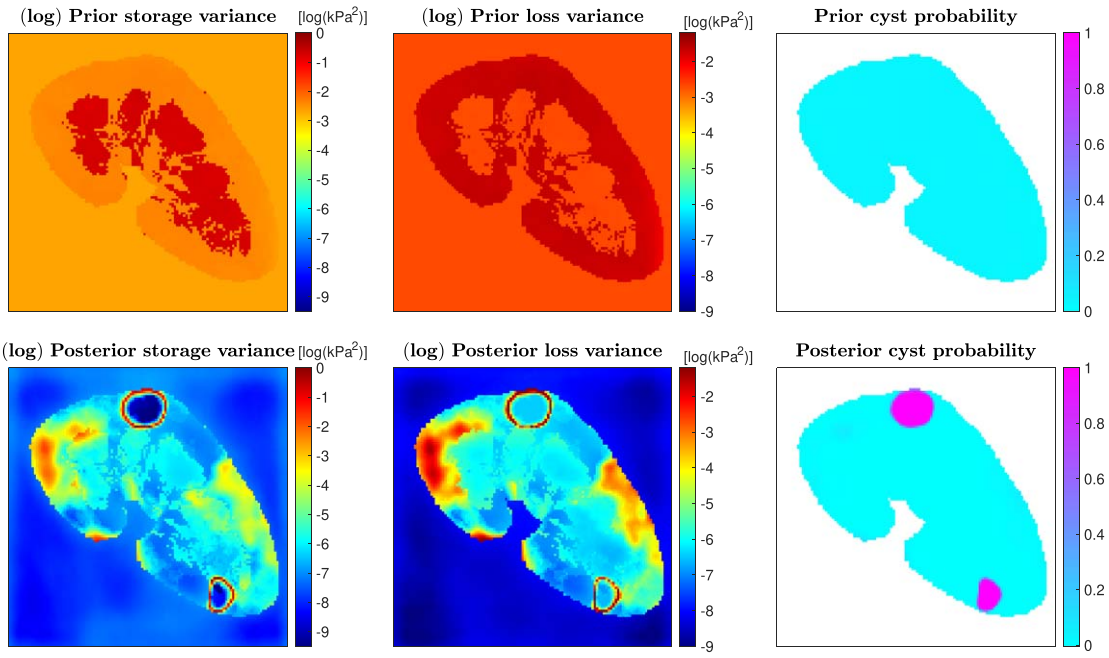

Figure 6. Top-left: prior (log) variance for the storage modulus [log (kPa^2)]. Top-middle: Prior (log) variance for the loss modulus [log (kPa^2)]. Bottom-left: posterior (log) variance for the storage modulus [log (kPa^2)]. Top-middle: posterior (log) variance for the loss modulus [log (kPa^2)]. Top-right: prior probability for the disease region. Top-right: posterior probability for the disease region.

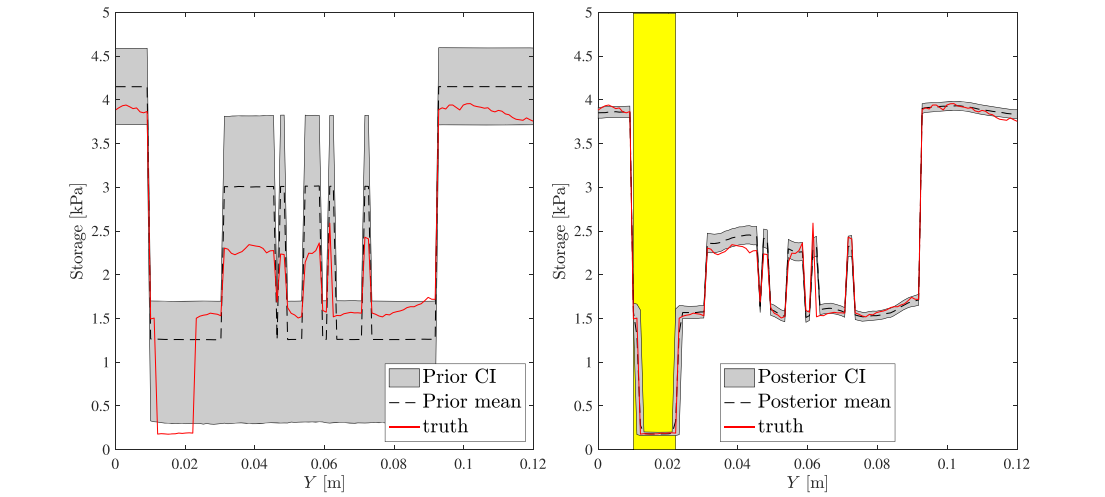

Figure 7. Left: Slice (at $y = 0.065$ m) of the graph of the true storage (red), prior mean (dotted black) and prior 95% credible intervals (shaded region). Right: slice (at $y = 0.065$ m) of the graph of the true storage (red), posterior mean (dotted black) and posterior 95% credible intervals (shaded region). The area highlighted in yellow corresponds to the location of the disease region.

algorithm. In figure 8 we plot the prior (dashed red) and posterior (dash-dotted blue) densities (here the posterior is approximated from samples via `ksdensity` in MATLAB) as well as the corresponding value of the ground truth (shown as a vertical solid line). We observe how from a uniform (uninformed) prior, the posterior produces Gaussian approximations concentrated around the truth. In most cases the truth is captured well within the posterior and in other cases the truth is on the tail of these posteriors.

In the *supplementary material* we study the effect of the ensemble size, J, as well as the level of measurements error/noise on the uncertainty measures that we compute using the EKI algorithm. In particular, we investigate the use of different choices of the parameters α_1 and α_2 in (32).

4.3. 3D experiments

In this section, we demonstrate the capabilities of EKI for inferring stiffness properties in the 3D forward model implemented as described in section 4.1. As for the 2D case, we assume that the kidney consists of cortex and

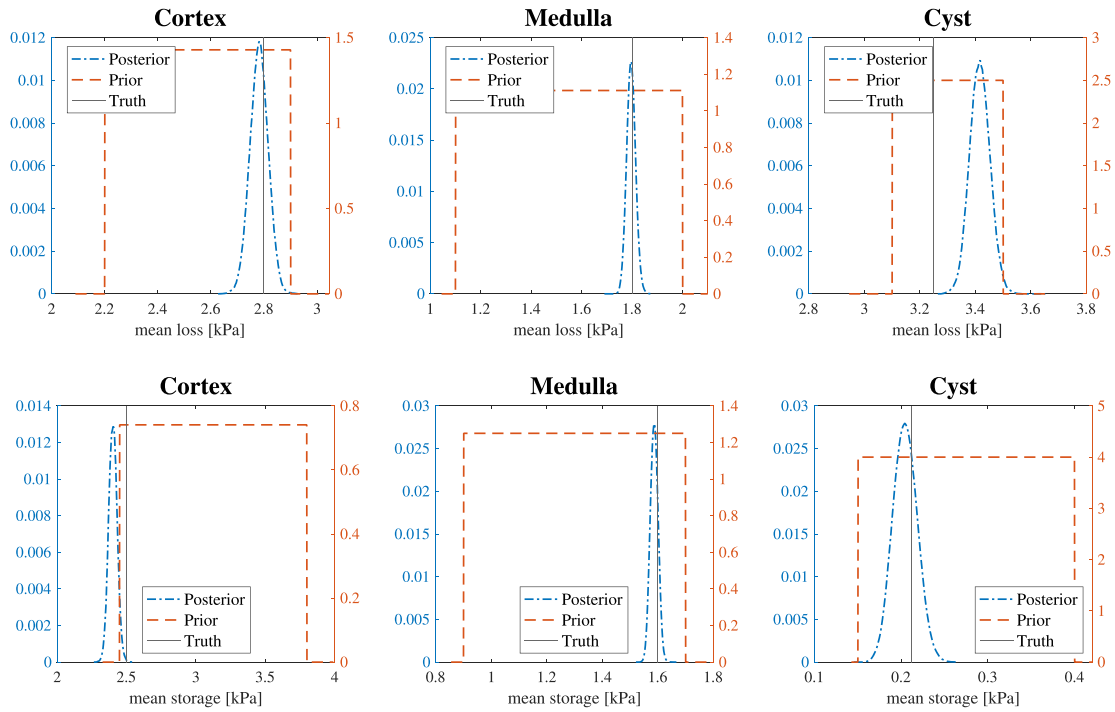

Figure 8. Prior (dashed red) and posterior (dash-dotted blue) densities (dimensionless) for the loss (top) and storage (bottom) modulus on each of the sub-regions of Ω. The corresponding value of the ground truth is shown as a vertical solid line.

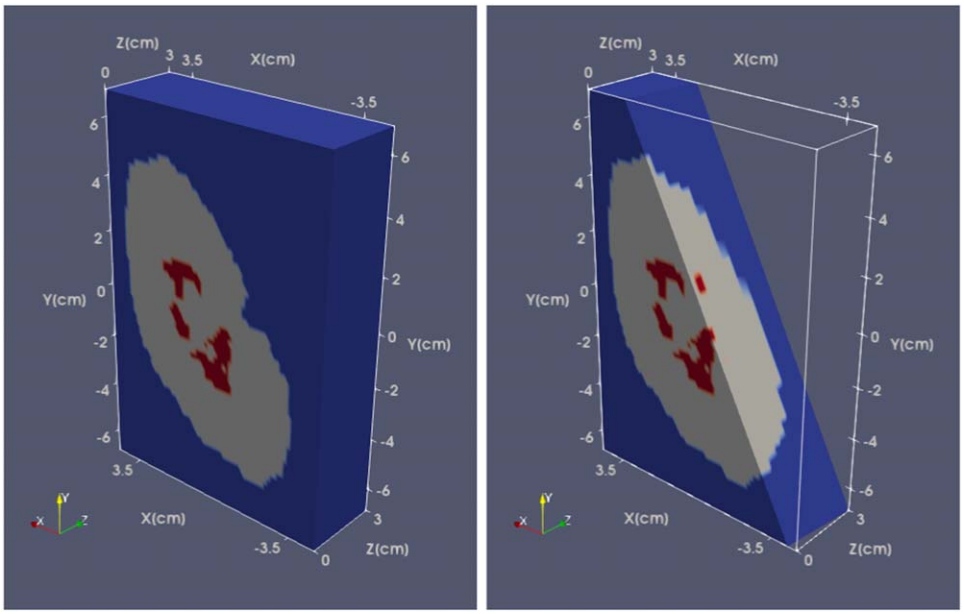

Figure 9. 3D Experiments. Setting for the known kidney geometry.

medulla that is immersed in a host/background media or tissue. We also assume that these regions are known and have the geometry shown in figure 9 which is obtained from an *in vivo* MR imaging of the kidney of a healthy volunteer, with signal thresholding performed to segment the medulla and cortex of the kidney. The ground truth for the storage and loss modulus are given on the top and middle-bottom panels of figure 10. These are, again, Gaussian RFs that we produce as described in section 3.1 with parameters specified in appendix B and implemented as discussed in appendix A.1. We have added two disconnected regions within the cortex of the kidney with anomalous values corresponding cystic lesions. Two views of these cystic regions are given in red at the top panels of figure 11. We have defined a sharp discontinuity in the storage values between the cyst and

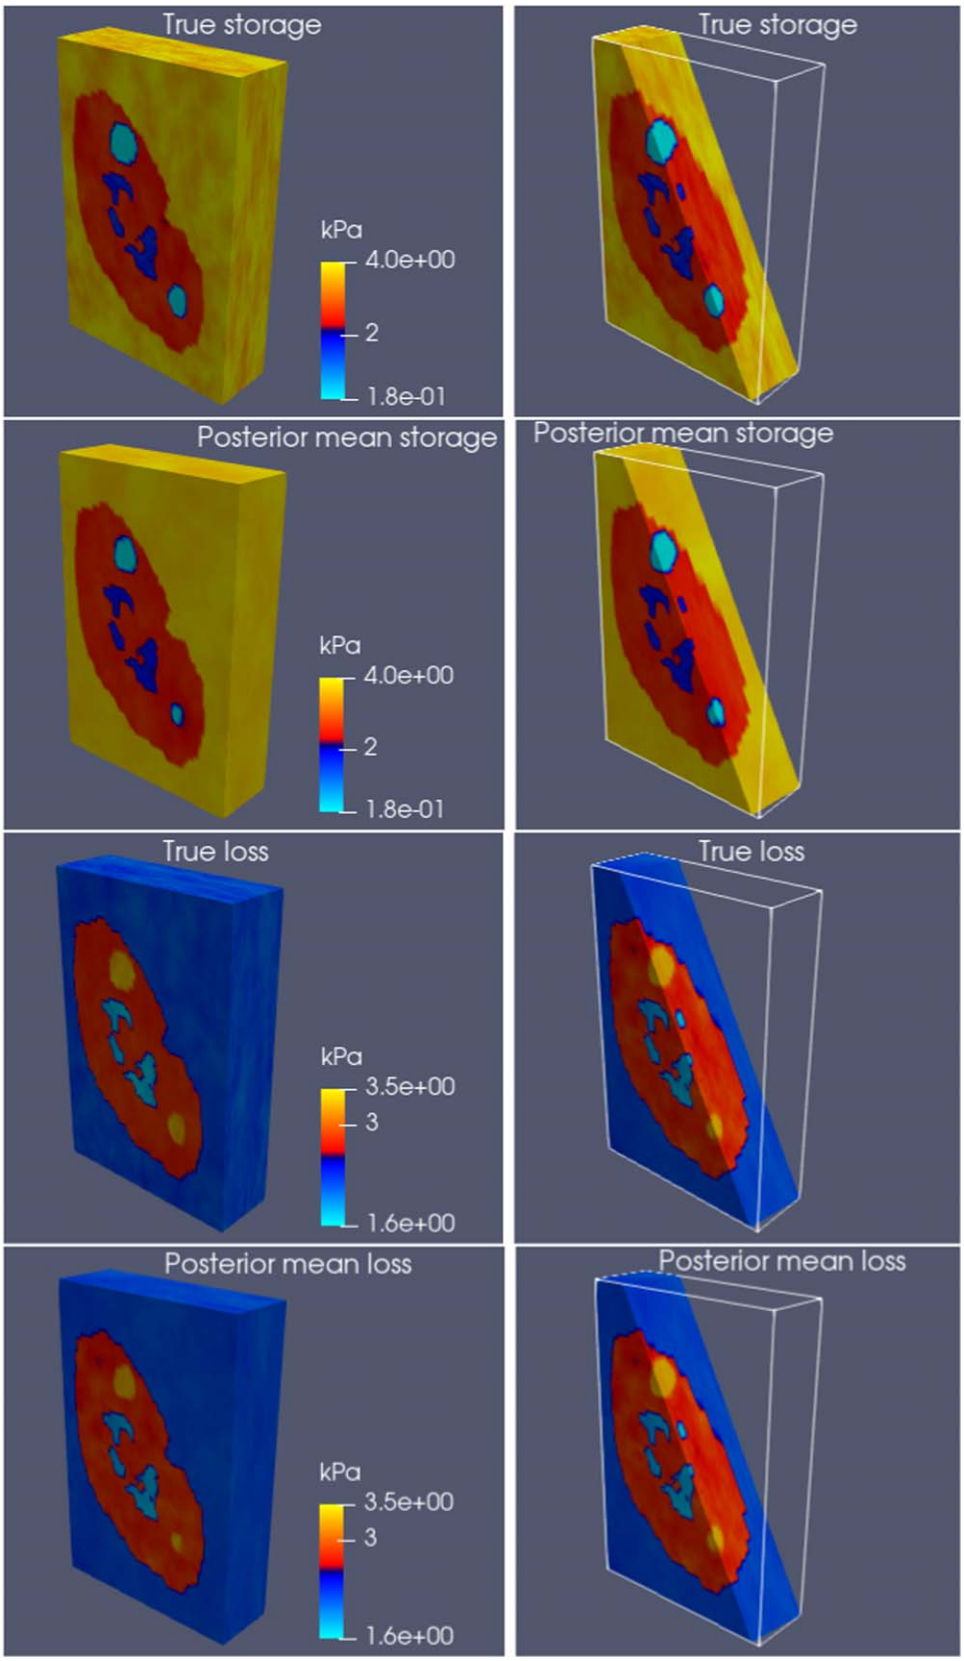

Figure 10. Top: true storage modulus [kPa]. Top-middle: posterior mean for the storage modulus [kPa]. Middle-bottom: true loss modulus [kPa]. Bottom: posterior mean for the loss modulus [kPa].

cortex regions. Note, the small dark blue region at the boundary of the cyst is simply an artifact from the interpolation employed for the visualisation tool.

We use the ground truth for the storage and loss, suitably interpolated on the `pdetool`-generated 3D mesh (tetrahedral elements), to produce displacements by solving equation (1) in the 3D domain. Note that for this

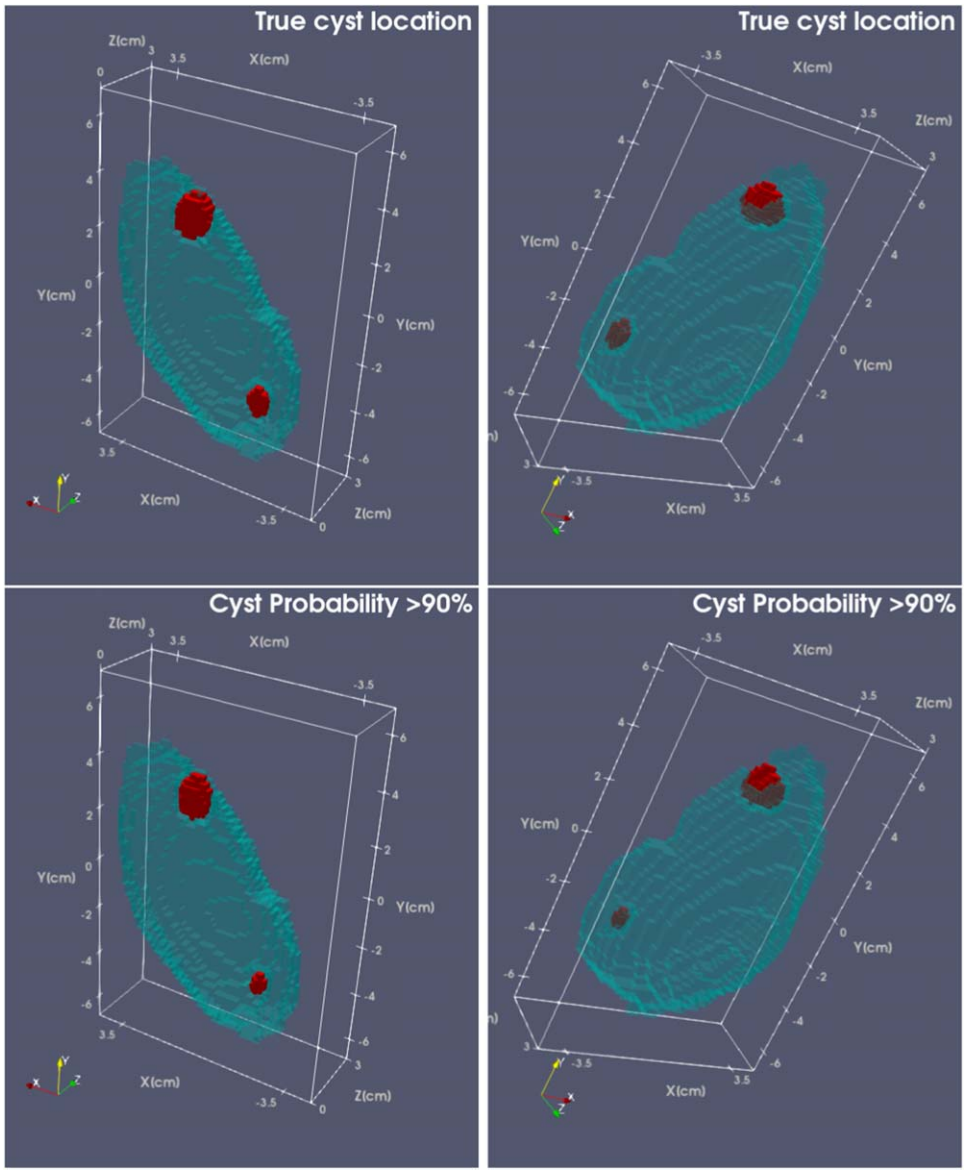

Figure 11. Top: views of the location of the true cystic region (in red). Bottom: the region (in red) for which the posterior probability within the cyst is greater than 90%.

case, the complex displacement field, $\mathbf{u}_M^\dagger$, has three components which are all used for the inversion. For simplicity, we use the nodes of the `pdetool` 3D mesh as the evaluation points, $\{\mathbf{x}_m\}_{m=1}^M$, for these displacements. Synthetic data $\mathbf{d}^\dagger$ are computed by adding Gaussian noise to the displacements as described for the 2D case. We consider only the case with relatively small measurement errors and so we employ $\alpha_1 = 0.025$ and $\alpha_2 = 5 \times 10^{-4}$ in (32). Since we use a non-uniform mesh with tetrahedral elements, we report the resolution of the synthetic MRE data in terms of the lengths of the shortest (2 mm) and longest (4 mm) edges in the mesh.

The choice of priors is done analogous to that of the 2D case (see also appendix B). Some prior samples for the storage are shown in the *supplementary material*, where, as for the 2D case, we note that our choice of prior yields a large degree of variability across samples, not only in terms of storage/loss modulus values but also in the shape and location of the disease region.

For the 3D case, the computational cost of each evaluation of the parameter-to-output map is substantially larger than for the 2D setting in the previous experiments. Even though we employ a relatively coarse finite element approximation, with approximately 47 000 elements compared to the 40 000 used in 2D, for solving the forward model (equation (1)), the parametrisation, $\mathcal{P}$, which involves solving additional PDEs (see equation (A1), from appendix A) requires a larger resolution that, in turn, increases the overall computational cost of EKI. As discussed in appendix A.1, we use 10^4 cells for the numerical approximation of the parameterisation in 2D while 147 000 cells are employed for the 3D case. Due to the increased computational cost of each run of the parameter-to-output map, we consider only an ensemble of $J = 2500$ particles. Our

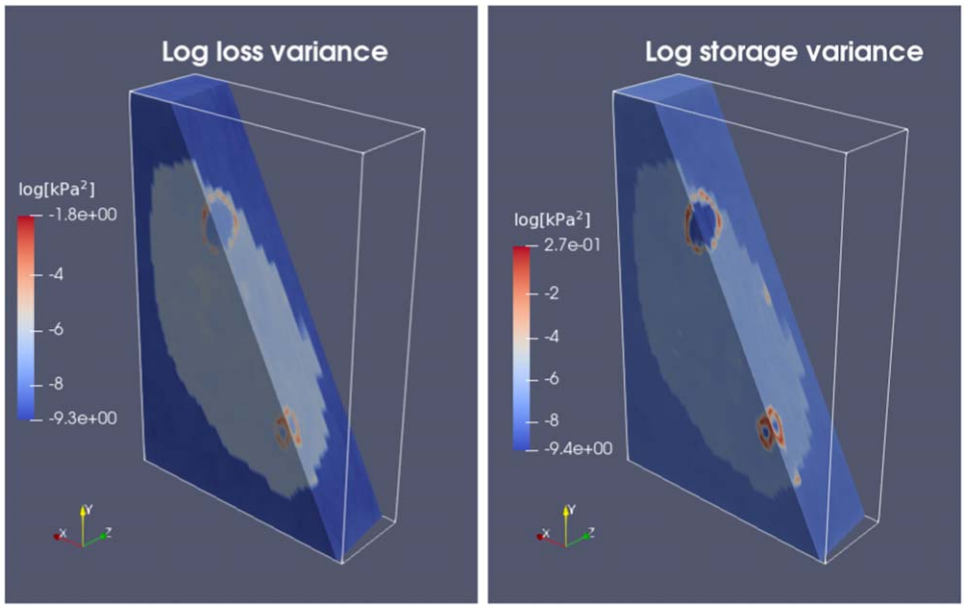

Figure 12. Log of the posterior variance for the loss (left) and storage (right) modulus.

experiments in section 4.2 (see also *supplementary material*) provided evidence that with such an ensemble size, the EKI algorithm produces robust and reliable posterior estimates of uncertainty.

The plots of the posterior means for storage and loss that we obtain via EKI are displayed at the top-middle and bottom panels of figure 10. The corresponding posterior variances are shown in figure 12. Similar to the 2D experiments we note that higher uncertainty (variance) is obtained at the boundary of the cystic region.

In the bottom panels of figure 11 we show the regions (in red) that corresponds to the volumes in which the posterior probability (computed via (28)) is greater than 90%. While the posterior mean shows a very good approximation to the values and variability of the storage/loss modulus, the disease region was also inferred accurately. It is quite remarkable how the EKI algorithm is able to capture the true location of the cyst within the region of high posterior probability (see bottom panels figure 11).

It is worth noticing that the variance of the small cyst is larger (see figure 12) than that of the small cyst in the 2D example (see figure 6). This explains why the size of the region of high probability corresponding to this cyst is smaller than the true (small) cyst (see figure 11). A possible reason for the larger uncertainty in the 3D case is that for each slice we employed, due to the computational burden of the 3D case, a lower resolution (70×70) compared to the resolution used for the 2D examples (100×100).

Despite the reconstructions results from this section expect the 2D results to be more accurate than the 3D case because each slice of the 3D set has a lower resolution than the single slice used in the 2D example. However, it appears that the posterior mean estimates of the cyst (see figure 10) seem to capture the size of the cyst very accurately and quite similar to those in the 2D case. But, there is larger uncertainty in the 3D estimates which explains why the region with probability greater than 90% is smaller than the true cyst.

5. Conclusions

A new and powerful method for MRE inversion has been introduced based on the EKI algorithm that uses RFs to accurately characterise spatially varying storage/loss modulus. In contrast to existing iterative approaches which often yield overly-smoothed elastograms, the EKI algorithm is equipped with a level-set parameterisation that allows the accurate inference of sharp interfaces between healthy and diseased tissue. Furthermore, EKI does not require differentiation of data which, due to the amplification of imaging-related noise, can lead to artifacts which are detrimental to the diagnostic capabilities of MRE. In addition, uncertainty of the reconstructed material properties can be quantified thanks to the probabilistic nature of the method. Since the forward viscoelastic model is encoded within the EKI algorithm only in a black box fashion, it can be easily modified to accommodate any material constitutive model (e.g. in a future work poroelastic materials could be explored). Also, the illustrative experiments shown here represent kidney MRE data but EKI could equally well be used for any other application such as brain or liver. Here the method has been shown to detect a renal cyst, but such uncertainty measures could be used by clinical practitioners to accurately compute and assess disease margins

such as for cancer, e.g. tumours in the kidney. Also, the accurate assessment of tumour margins is important for many organs, e.g. prostate, in terms of prognosis and determining the likelihood of secondary disease outside of the primary site, and for brain tumours where the precise location of disease margins is important for surgery planning. While prior information on anatomical zones of the cortex and medulla as informed by non-elastography-based MRI data (e.g. T1- or T2-weighted imaging) is included in this methodology, the methodology can also be used in organs without such anatomical definitions. For example, the liver does not have such anatomical boundaries, but liver fibrosis can be found by this method to high accuracy and without prior information. The precise location and extent of liver fibrosis are important when staging the progression of disease or therapy response, and also for biopsy planning.

While the work presented demonstrates an initial exploration of the methodology with simulated data, it convincingly represents the ready applicability of the technique to real acquired 3D MRE data. The application of EKI with acquired data would require the use of measured displacements as boundary conditions for the forward model which here we prescribed for the sake of illustration. However, when boundary conditions are also subject to noise, the EKI framework can be extended to incorporate these uncertainties within the inversion algorithm (Iglesias *et al* 2018). Finally, additional sources of modeling errors are likely to arise when dealing with realistic settings. Fortunately, the Bayesian setting in which the EKI algorithm is derived, enable us to account, and possibly infer those errors within the inversion methodology (Calvetti *et al* 2018).

Further validation of the proposed method should be performed with real data obtained under control conditions, i.e. where a ground truth can be well defined (e.g. using gel phantoms). Such experiments will allow us to compare outcomes of the EKI algorithm obtained based on real data and on simulated data related to the used experimental set-up. This comparison would enable us to establish (and possibly incorporate within the inversion) additional sources of uncertainties arising from inherent errors in the models that are used to describe MRE data. Experiments with real acquired data are also needed to further validate, potentially improve or even reformulate our modeling assumptions. Indeed, recall that our inversion framework comprises numerous assumptions on the (i) the forward (visco-elastic) model, (ii) the random field characterisation of stiffness properties combined with the level-set parameterisation of diseased tissue, and (iii) the priors that we have selected for the parameters within those parameterisations. If it will be required, alternative characterisations of stiffness properties and/or other type of edge-preserving (e.g. total variation) priors could be used within the Bayesian setting. Having real data from controlled experiments would enable us to compare, e.g. via Bayesian model comparison techniques (Gelman *et al* 2004), the modeling assumptions that best explain the data and thus the modeling settings that provide us with an accurate MRE characterisation of the elastic properties of biological tissue.

Acknowledgments

This work was supported by a Kidney Research UK Innovation Grant KS_RP_016_20190919. This research was also funded by the NIHR Nottingham Biomedical Research Centre and carried out at/supported by the NIHR Nottingham Clinical Research Facilities. The views expressed are those of the author(s) and not necessarily those of the NHS, the NIHR or the Department of Health and Social Care. We are grateful for access to the University of Nottingham's Augusta HPC service. The first author would like to thank Daniele Avitabile for very useful discussions on this topic.

Ethical statement

The 2D data of this study was derived from *ex vivo* kidney MRI data. The 3D data was derived from the *in vivo* MRI scan of a healthy subject as part of a study which was also approved by the University of Nottingham Medical School Ethics Committee (Approval number: E14032013). This subject signed a consent form giving us permission to publish data. All research was conducted in accordance with the principles embodied in the Declaration of Helsinki and in accordance with local statutory requirements.

Appendix A. Whittle–Matern parameterisations

In this appendix we introduce the mapping $\mathcal{P}^{WM}$ that we use in (19) to parameterise the level-set function f (that determines the cyst region), as well as the storage and loss modulus, $G_{s,\beta}$ and $G_{l,\beta}$ ($\beta \in \mathcal{B} = \{b,\ c,\ m,\ t\}$), for each constituent of the geometry under consideration.

Let us consider first the 3D case ($d = 3$). For fixed $\alpha \in \{s, l\}$ and $\beta \in \mathcal{B}$, we define the function (random field) $\log G_{\alpha,\beta}(\mathbf{x})$ as a solution of the fractional stochastic PDE:

$$[I - \nabla \cdot \mathrm{diag}(L_x^2, L_y^2, L_z^2)\nabla]^{\zeta/2}(\log G_{\alpha,\beta}(\mathbf{x}) - \log m) = \left[2^d L_x L_y L_z \sigma^2 \pi^{d/2} \frac{\Gamma(\zeta)}{\Gamma(\zeta - d/2)}\right]^{1/2} \omega(\mathbf{x}), \quad \mathbf{x} \in \Omega, \tag{A1}$$

where I is the identity operator, Γ is the gamma function and $\omega(\mathbf{x})$ is a (generalised) random function. Furthermore, ζ is a parameter that controls the smoothness of the function $\log G_{\alpha,\beta}$, σ is an amplitude scale, $\log m$ is the mean value of $\log G_{\alpha,\beta}$, and L_x, L_y and L_z are (positive) numbers that describe intrinsic length-scales along the x, y and z directions, respectively. In (A1) $\mathrm{diag}(L_x^2, L_y^2, L_z^2)$ denotes a 3×3 diagonal matrix with the vector (L_x^2, L_y^2, L_z^2) on its diagonal.

For the case $d = 2$, in equation (A1) we replace the matrix $\mathrm{diag}(L_x^2, L_y^2, L_z^2)$ with $\mathrm{diag}(L_x^2, L_y^2)$ and the product $L_x L_y L_z$ with $L_x L_y$.

After we define appropriate boundary conditions for (A1), we can solve this PDE problem to find $\log G_{\alpha,\beta}(\mathbf{x})$ provided the following input parameters are specified:

$$\Theta_{\alpha,\beta} \equiv (m, \sigma, \zeta, L_x, L_y, L_z, \omega). \tag{A2}$$

In other words, equation (A1) (together with suitable boundary conditions) defines the mapping $\mathcal{P}^{WM}$:

$$\Theta_{\alpha,\beta} \mapsto \mathcal{P}^{WM}(\Theta_{\alpha,\beta}) = \log G_{\alpha,\beta}. \tag{A3}$$

Note that all the terms in the right-hand side of (A2) should be labeled to reflect the particular choice of α and β. While such a labeling is omitted for clarity in the notation, we shall understand that every function $\log G_{\alpha,\beta}(\mathbf{x})$ is parameterised by a unique set of inputs comprised in $\Theta_{\alpha,\beta}$. The selection of these parameters is discussed in appendix B.

The motivation for using the parameterisation (A3) stems from the work of Lindgren *et al* (2011). They showed that if $\omega(\mathbf{x})$ is a Gaussian white noise, then the function $\log G_{\alpha,\beta}(\mathbf{x})$ is a Gaussian RF with mean $\log m$ and covariance operator with Matern auto-correlation function (Matérn 1986, Stein 1999, Lasanen *et al* 2014):

$$\mathrm{ACF}(\mathbf{x}) = \sigma^2 \frac{1}{2^{\nu-1}\Gamma(\nu)} \|\mathbf{x}\|_L^{\nu} K_\nu(\|\mathbf{x}\|_L), \tag{A4}$$

where $\nu = \zeta - d/2$, K_ν denotes the modified Bessel function of the second kind of order ν, and

$$\|\mathbf{x}\|_L \equiv \sqrt{\frac{x^2}{L_x^2} + \frac{y^2}{L_y^2} + \frac{z^2}{L_z^2}}.$$

Alternative methods to the SPDE formulation in (A1) can be used to represent Gaussian RFs. For example, one can represent such fields in terms of their Karhunen–Loeve expansion. However, the advantage of our formulation is that it enables us to easily include hyperparameters, such as the intrinsic length scales, within the EKI framework. Inferring these parameters have been shown to be essential for the accurate inference of physical properties within the Bayesian framework (Dunlop *et al* 2017, Matveev *et al* 2021).

Although we introduce the parameterisation in terms of the functions $\log G_{\alpha,\beta}$, the same parameterisation is employed for the level-set function f as indicated in (19).

A.1. Numerical implementation

We consider only the case in which the smoothness parameter, ζ, is a positive even integer so that (A1) can be computed using standard numerical methods as proposed in Lindgren *et al* (2011). To this end, we use a bespoke MATLAB implementation of the cell-centered finite difference method (Russell and Wheeler 1983). We employ regular grids of 100×100 and $70 \times 70 \times 30$ cells for the examples in 2D and 3D, respectively. We use Robin boundary conditions which have been shown in Lasanen *et al* (2014) to be optimal in terms of reducing boundary effects. For a given set of inputs $\boldsymbol{\Theta}$, we use our numerical implementation of (A1) to compute numerical approximations of functions $\log G_{\alpha,\beta}$'s and f (i.e. this corresponds to the numerical approximation of $\mathcal{P}^{WM}(\boldsymbol{\Theta})$). The next step is to compute the storage/loss modulus over the entire domain Ω. This is done via the mapping $\mathcal{Q}$ (see equation (18)) which is defined via equation (16). Even though our implementation of (A1) computes each $G_{\alpha,\beta}$ over the entire domain Ω, we recall that only the values of $G_{\alpha,\beta}$ on the corresponding region of the tissue are utilised (e.g. $G_{s,c}$ is used only for the cortex). Once the storage and loss modulus $(G_s, G_l) = \mathcal{P}(\boldsymbol{\Theta})$ have been numerically approximated (for a given $\boldsymbol{\Theta}$) on our regular grid in which (A1) was solved, we proceed to interpolate these functions on the finite element grid that we then use for solving (1) as discussed in section 4.1.

A.2. Fixed parameters

For simplicity, for each choice of α and β we keep the amplitude scale, σ, and the smoothness parameter, ζ, constant. This means that we leave these out of the inference scheme while we focus on inferring length scales L_x, L_y, L_z as well as $\omega(\mathbf{x})$ and m. The values that we use for σ corresponding to the $\boldsymbol{\Theta}_{\alpha,\beta}$'s are shown in table B1. These

σ's are selected so that the values of the (prior) random samples of storage and loss modulus (obtained by exponentiation of (A3)) belong to a reasonable and physically-consistent range. Accordingly, we choose the smoothness parameter ζ so that these samples exhibit a realistic amount of spatial variability. We find that the value $\zeta = 2$ results in samples of reasonably smooth (e.g. at least continuous) functions (see top panels of figure 2). This level of smoothness is suitable to reflect the fact that, within the same tissue type, the stiffness properties are not expected to drastically change unless there is a cyst/anomaly (which we characterise as a separate region via the level-set function). For the level set function, f, we choose $\sigma = 1$ and $\zeta = 4$ which, together with our selection of threshold $\kappa = 1.75$, yield κ-levels that provide (visually) well-defined interfaces between normal and diseased regions (see middle and right panels of figure 1).

Appendix B. The prior

As discussed in section 3.2, the proposed EKI approach for MRE consists of inferring, within the Bayesian framework, the parameters $\boldsymbol{\Theta}_{\alpha,\beta}$ ($\alpha \in \{s, l\}$, $\beta \in \mathcal{B}$) and Θ_f which characterise, via the parameterisation $\mathcal{P}^{WM}$ introduced in appendix A (see (19)), the storage/loss modulus as well as the level-set function that determines the cystic region. We also recall that within the Bayesian framework we need to specify a prior distribution, $\mathbb{P}(\boldsymbol{\Theta})$, on those parameters which are comprised in $\boldsymbol{\Theta}$ (see equation (20)). This prior is crucial for EKI as the initial ensemble in algorithm 1 consists of samples from this prior.

Let us assume that, under the prior, all the parameters comprised in $\boldsymbol{\Theta}$ are independent so that

$$\mathbb{P}(\boldsymbol{\Theta}) = \prod_{\beta\in\mathcal{B}} \mathbb{P}(\boldsymbol{\Theta}_{s,\beta}) \prod_{\beta\in\mathcal{B}} \mathbb{P}(\boldsymbol{\Theta}_{l,\beta})\mathbb{P}(\boldsymbol{\Theta}_s). \tag{B1}$$

We recall that each of the $\boldsymbol{\Theta}_{\alpha,\beta}$'s as well as $\boldsymbol{\Theta}_s$ have the form given in (A2) but excluding σ and ζ which are fixed (so we leave them out from the inference algorithm). For the sake of exposition, we focus only on $\boldsymbol{\Theta}_{\alpha,\beta}$ and define $\boldsymbol{\Theta}_f$ by analogy. We use the following assumption of independence:

$$\mathbb{P}(\boldsymbol{\Theta}_{\alpha,\beta}) = \mathbb{P}(m)\mathbb{P}(L_x)\mathbb{P}(L_y)\mathbb{P}(L_z)\mathbb{P}(\omega), \tag{B2}$$

where $\mathbb{P}(m)$, $\mathbb{P}(L_x)$,..., denote the priors of m, L_x,..., and where, again, for clarity in the notation we have omitted the dependence on (α, β) for each of the priors in the right hand side of (B2). We now proceed to define each of these priors.

For each of our parameters $\boldsymbol{\Theta}_{\alpha,\beta}$ and $\boldsymbol{\Theta}_f$ the corresponding ω is chosen to be Gaussian white noise (we denote this by $\mathbb{P}(\omega) = N(0, I)$). We reiterate that with this choice of ω, the work of Lindgren *et al* (2011) ensures that the Whittle–Matern parameterisation from appendix A yields samples $\log G_{\alpha,\beta}$ and f that are Gaussian with Matern covariance and mean $\log m$. This implies that the $G_{\alpha,\beta}$ are log-normal RFs with mean m. The prior distributions for the mean m associated with the $G_{\alpha,\beta}$'s are displayed in table B1. With this selection of priors on m we ensure that the average values for the prior random samples of loss/storage modulus belong to an acceptable range that is consistent with the values of biological tissue corresponding to each anatomical regions. For the sake of demonstration, here we have used wide (uninformative) uniform priors. In practice, however, we can use any available knowledge (i.e. from previous studies or literature) to design other type of prior distribution (not necessarily uniform) that better aligns with our prior knowledge of the anatomical tissue under consideration.

For all choices of (α, β) we select the same prior distribution for the length scales L_x, L_y and L_z. We consider the following

$$\begin{aligned} \mathbb{P}(L_x) = U[D_x/20, D_x/4], \quad \mathbb{P}(L_y) = U[D_y/20, D_y/4], \\ \mathbb{P}(L_z) = U[D_z/20, D_z/4], \end{aligned} \tag{B3}$$

where D_x, D_y and D_z are the lengths of the edges/faces of the rectangular/tetrahedral domain. This selection is motivated by the known fact (e.g. see Lasanen *et al* 2014) that values of a sample of a Gaussian RFs (with Matern covariance given by (A4)) have a correlation of approximately 0.1 at a distance of $L_x\sqrt{8(\zeta - d/2)}$ in the x-direction. With the prior from (B3), such a distance (in the 2D case with our choice $\zeta = 2$) is between 0.14 and 0.7. This variability of the prior length scales will allow us to characterise stiffness properties that can vary more quickly in one particular direction. These can be observed in the top-middle row panels of figure 6 in the *supplementary material*. Indeed, for some of those prior samples of the loss modulus, the background region displays very long correlations in only one direction.

Table B1. Relevant prior information. $U[a, b]$ denotes the uniform distribution on the interval $[a, b]$.

α	β	σ [log (Pa)]	$\mathbb{P}(m)$
s	b	0.01	U[3.7 kPa, 4.6 kPa]
s	c	0.03	U[0.9 kPa, 1.7 kPa]
s	m	0.03	U[2.45 kPa, 3.8 kPa]
s	d	0.1	U[0.15 kPa, 0.4 kPa]
l	b	0.01	U[2.0 kPa, 2.9 kPa]
l	c	0.03	U[1.1 kPa, 2.0 kPa]
l	m	0.03	U[2.2 kPa, 2.9 kPa]
l	d	0.01	U[3.1 kPa, 3.5 kPa]

ORCID iDs

Marco Iglesias https://orcid.org/0000-0002-8952-717X

References

Alberty J, Carstensen C, Funken S A and Klose R 2002 Matlab implementation of the finite element method in elasticity *Computing* **69** 239–63
Ammari H, Seo J K and Zhou L 2015 Viscoelastic modulus reconstruction using time harmonic vibrations *Math. Model. Anal.* **20** 836–51
Calvetti D, Dunlop M, Somersalo E and Stuart A 2018 Iterative updating of model error for Bayesian inversion *Inverse Prob.* **34** 025008
Chada N K, Iglesias M A, Roininen L and Stuart A M 2018 Parameterizations for ensemble Kalman inversion *Inverse Prob.* **34** 055009
Ciarlet P G 1988 *Mathematical Elasticity. Vol.1 Three Dimensional Elasticity—Mathematical Elasticity* (Amsterdam: Elsevier)
Doyley M M 2012 Model-based elastography: a survey of approaches to the inverse elasticity problem *Phys. Med. Biol.* **57** R35–73
Dunlop M M, Helin T and Stuart A M 2020 Hyperparameter estimation in Bayesian MAP estimation: parameterizations and consistency *SMAI J. Comput. Math.* **6** 69–100
Dunlop M M, Iglesias M A and Stuart A M 2017 Hierarchical Bayesian level set inversion *Stat. Comput.* **27** 1555–84
Engl H, Hanke M and Neubauer A 1996 *Regularization of Inverse Problems* vol 375 (Berlin: Springer)
Fovargue D, Nordsletten D and Sinkus R 2018 Stiffness reconstruction methods for MR elastography *NMR Biomed.* **31** e3935
Fovargue D, Nordsletten D and Sinkus R 2018 Stiffness reconstruction methods for MR elastography *NMR Biomed.* **31**
Gelman A, Carlin J B, Stern H S and Rubin D B 2004 *Bayesian Data Analysis* 2nd edn (Boca Raton, FL: CRC Press)
Higashimori N 2007 Identification of viscoelastic properties by magnetic resonance elastography *J. Physics: Confer. Ser.* **73** 012009
Hiscox L V, Johnson C L, Barnhill E, McGarry M D, Huston J, van Beek E J, Starr J M and Roberts N 2016 Magnetic resonance elastography (MRE) of the human brain: technique, findings and clinical applications *Phys. Med. Biol.* **61** R401–37
Hu L and Shan X 2020 Enhanced complex local frequency elastography method for tumor viscoelastic shear modulus reconstruction *Comput. Methods Programs Biomed.* **195** 105605
Iglesias M, Lu Y and Stuart A M 2016 A Bayesian level set method for geometric inverse problems *Interfaces and Free Boundaries* **18** 181
Iglesias M, Park M and Tretyakov M V 2018 Bayesian inversion in resin transfer molding *Inverse Prob.* **34** 105002
Iglesias M, Sawlan Z, Scavino M, Tempone R and Wood C 2018 Ensemble-marginalized Kalman filter for linear time-dependent PDEs with noisy boundary conditions: application to heat transfer in building walls *Inverse Prob.* **34** 075008
Iglesias M and Yang Y 2021 Adaptive regularisation for ensemble Kalman inversion *Inverse Prob.* **37** 025008
Iglesias M A 2014 Iterative regularization for ensemble data assimilation in reservoir models *Comput. Geosci.* **19** 177–212
Iglesias M A 2016 A regularizing iterative ensemble Kalman method for PDE-constrained inverse problems *Inverse Prob.* **32** 025002
Iglesias M A, Law K J H and Stuart A M 2013 The ensemble Kalman filter for inverse problems *Inverse Prob.* **29** 045001
Jiang Y, Fujiwara H and Nakamura G 2011 Approximate steady state models for magnetic resonance elastography *SIAM J. Appl. Math.* **71** 1965–89
Jiang Y, Fujiwara H and Nakamura G 2020 Erratum for 'approximate steady state models for magnetic resonance elastography' *SIAM J. Appl. Math.* **80** 2001
Jiang Y and Qian S-h 2020 Bayesian approach for recovering piecewise constant viscoelasticity from MRE data *Acta Math. Appl. Sinica, English Ser.* **36** 223–36
Kaipio J and Somersalo E 2005 *Statistical and Computational Inverse Problems* (Dordrecht: Springer)
Kennedy P, Macgregor L J, Barnhill E, Johnson C L, Perrins M, Hunter A, Brown C, van Beek E J R and Roberts N 2017 MR elastography measurement of the effect of passive warmup prior to eccentric exercise on thigh muscle mechanical properties *J. Magn. Reson. Imaging* **46** 1115–27
Landau L D and Lifschitz E M 1986 *Theory of Elasticity* (Oxford: Pergamon)
Lasanen S, Huttunen J M J and Roininen L 2014 Whittle–Matérn priors for Bayesian statistical inversion with applications in electrical impedance tomography *Inverse Probl. Imaging* **8** 561–86
Li B N, Chui C K, Ong S H, Numano T, Washio T, Homma K, Chang S, Venkatesh S and Kobayashi E 2012 Modeling shear modulus distribution in magnetic resonance elastography with piecewise constant level sets *Magn. Reson. Imaging* **30** 390–401
Li H, Flé G, Bhatt M, Qu Z, Ghazavi S, Yazdani L, Bosio G, Rafati I and Cloutier G 2021 Viscoelasticity imaging of biological tissues and single cells using shear wave propagation *Front. Phys.* **9** 666192
Li S, Chen M, Wang W, Zhao W, Wang J, Zhao X and Zhou C 2011 A feasibility study of MR elastography in the diagnosis of prostate cancer at 3.0T *Acta Radiol.* **53** 354–8
Lindgren F, Havard R and Lindström J 2011 An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach *J. R. Statist. Soc.* B **73** 423–98

Manduca A, Muthupillai R, Rossman P J, Greenleaf J F and Ehman R L 1996 Visualization of tissue elasticity by magnetic resonance elastography *Visualization in Biomedical Computing* ed K H Höhne and R Kikinis (Berlin, Heidelberg: Springer Berlin Heidelberg) pp 63–8

Manduca A, Oliphant T E, Dresner M A, Mahowald J L, Kruse S A, Amromin E, Felmlee J P, Greenleaf J F and Ehman R L 2001 Magnetic resonance elastography: non-invasive mapping of tissue elasticity *Med. Image Anal.* **5** 237–54

Marinelli J P, Levin D L, Vassallo R, Carter R E, Hubmayr R D, Ehman R L and McGee K P 2017 Quantitative assessment of lung stiffness in patients with interstitial lung disease using MR elastography *J. Magn. Reson. Imaging* **46** 365–74

Matérn B 1986 *Spatial Variation, Lecture Notes in Statistics, No. 36* (Berlin: Springer)

Mathew R P and Venkatesh S K 2018 Imaging of hepatic fibrosis *Curr. Gastroenterol. Rep.* **20** 45

Matveev M Y, Endruweit A, Long A C, Iglesias M A and Tretyakov M V 2021 Bayesian inversion algorithm for estimating local variations in permeability and porosity of reinforcements using experimental data *Composites* A **143** 106323

McGarry M, Johnson C L, Sutton B P, Van Houten E E W, Georgiadis J G, Weaver J B and Paulsen K D 2013 Including spatial information in nonlinear inversion MR elastography using soft prior regularization *IEEE Trans. Med. Imaging* **32** 1901–9

McGarry M, Van Houten E E W, Solamen L, Gordon-Wylie S, Weaver J and Paulsen K D 2019 Uniqueness of poroelastic and viscoelastic nonlinear inversion MR elastography at low frequencies *Phys. Med. Biol.* **64** 075006

McGrath D M 2018 Magnetic resonance elastography *Biomechanics of Soft Tissues: Principles and Applications* ed A Al Mayah (Boca Raton: CRC Press) pp 55–94

McGrath D M, Foltz W D, Al-Mayah A, Niu C J and Brock K K 2012 Quasi-static magnetic resonance elastography at 7 T to measure the effect of pathology before and after fixation on tissue biomechanical properties *Magn. Reson. Med.* **68** 152–65

Mohammadi N, Doyley M M and Cetin M 2021 Finite element reconstruction of stiffness images in MR elastography using statistical physical forward modeling and proximal optimization methods *2021 IEEE 18th Int. Symp. on Biomedical Imaging (ISBI)* pp 200–3

Muthupillai R, Lomas D J, Rossman P J, Greenleaf J F, Manduca A and Ehman R L 1995 Magnetic resonance elastography by direct visualization of propagating acoustic strain waves *Science* **269** 1854–7

Napoli M E, Goswami S, McAleavey S A, Doyley M M and Howard T M 2021 Enabling quantitative robot-assisted compressional elastography via the extended Kalman filter *Phys. Med. Biol.* **66** 225014

Papazoglou S, Hamhaber U, Braun J and Sack I 2008 Algebraic Helmholtz inversion in planar magnetic resonance elastography *Phys. Med. Biol.* **53** 3147–58

Russell T and Wheeler M 1983 Finite element and finite difference methods for continuous flows in porous media *The Mathematics of Reservoir Simulation* ed R E Ewing (Philadelphia, PA: SIAM) pp 35–106

Sakai N, Takehara Y, Yamashita S, Ohishi N, Kawaji H, Sameshima T, Baba S, Sakahara H and Namba H 2016 Shear stiffness of 4 common intracranial tumors measured using MR elastography: Comparison with intraoperative consistency grading *Am. J. Neuroradiol.* **37** 1851–9

Simon L D, Iglesias M, Jones B and Wood C 2018 Quantifying uncertainty in thermophysical properties of walls by means of bayesian inversion *Energy Build.* **177** 220–45

Singh S *et al* 2015 Diagnostic performance of magnetic resonance elastography in staging liver fibrosis: a systematic review and meta-analysis of individual participant data *Clin. Gastroenterol. Hepatol.* **13** 440–51

Sinkus R, Lorenzen J, Schrader D, Lorenzen M, Dargatz M and Holz D 2000 High-resolution tensor MR elastography for breast tumour detection *Phys. Med. Biol.* **45** 1649–64

Sinkus R, Tanter M, Xydeas T, Catheline S, Bercoff J and Fink M 2005 Viscoelastic shear properties of *in vivo* breast lesions measured by MR elastography *Magn. Reson. Imaging* **23** 159–65

Stein M L 1999 *Interpolation of Spatial Data: Some Theory for Kriging* (Berlin: Springer)

Stuart A 2010 Inverse problems: a Bayesian perspective *Acta Numer.* **19** 451–559

Tso C-H M, Iglesias M, Wilkinson P, Kuras O, Chambers J and Binley A 2021 Efficient multiscale imaging of subsurface resistivity with uncertainty quantification using ensemble Kalman inversion *Geophys. J. Int.* **225** 887–905

Van Houten E E W, Paulsen K D, Miga M I, Kennedy F E and Weaver J B 1999 An overlapping subzone technique for MR-based elastic property reconstruction *Magn. Reson. Med.* **42** 779–86

Venkatesh S K, Yin M and Ehman R L 2013 Magnetic resonance elastography of liver: technique, analysis, and clinical applications *J. Magn. Reson. Imaging* **37** 544–55

Venkatesh S K, Yin M, Glockner J F, Takahashi N, Araoz P A, Talwalkar J A and Ehman R L 2008 MR elastography of liver tumors: preliminary results *Amer. J. Roent.* **190** 1534–40

Yin M, Talwalkar J A, Glaser K J, Manduca A, Grimm R C, Rossman P J, Fidler J L and Ehman R L 2007 Assessment of hepatic fibrosis with magnetic resonance elastography *Clin. Gastroenterol. Hepatol.* **5** 1207–13

Zhang Y, Oberai A A, Barbone P E and Harari I 2012 Solution of the time-harmonic viscoelastic inverse problem with interior data in two dimensions *Int. J. Numer. Methods Eng.* **92** 1100–16