

Inference of Heterogeneous Material Properties via Infinite-Dimensional Integrated DIC

Joseph Kirchhoff^{a,*}, Dingcheng Luo^b, Thomas O’Leary-Roseberry^b, Omar Ghattas^{a,b}

^aWalker Department of Mechanical Engineering, The University of Texas at Austin, 204 E. Dean Keeton Street, C2200, Austin, TX 78712, USA

^bOden Institute for Computational Engineering and Sciences, The University of Texas at Austin, 201 E. 24th Street, C0200, Austin, TX 78712, USA

Abstract

We present a scalable and efficient framework for the inference of spatially-varying parameters of continuum materials from image observations of their deformations. Our goal is the nondestructive identification of arbitrary damage, defects, anomalies and inclusions without knowledge of their morphology or strength. Since these effects cannot be directly observed, we pose their identification as an inverse problem. Our approach builds on integrated digital image correlation (IDIC, Besnard Hild, Roux, 2006), which poses the image registration and material inference as a monolithic inverse problem, thereby enforcing physical consistency of the image registration using the governing PDE. Existing work on IDIC has focused on low-dimensional parameterizations of materials. In order to accommodate the inference of heterogeneous material property fields that are formally infinite dimensional, we present ∞ -IDIC, a general formulation of the PDE-constrained coupled image registration and inversion posed directly in the function space setting. This leads to several mathematical and algorithmic challenges arising from the ill-posedness and high dimensionality of the inverse problem. To address ill-posedness, we consider various regularization schemes, namely $H^1(\Omega)$ and total variation for the inference of smooth and sharp features, respectively. To address the computational costs associated with the discretized problem, we use an efficient inexact-Newton CG framework for solving the regularized inverse problem. In numerical experiments, we demonstrate the ability of ∞ -IDIC to characterize complex, spatially varying Lamé parameter fields of linear elastic and hyperelastic materials. Our method exhibits (i) the ability to recover fine-scale and sharp material features, (ii) mesh-independent convergence performance and hyperparameter selection, (iii) robustness to observational noise.

Keywords: Digital image correlation (DIC), inverse problems, nondestructive evaluation (NDE), material identification, dimension independence.

Contents

1	Introduction	2
1.1	Related works	3
1.2	Contributions	3
1.3	Layout of the paper	4
2	The ∞-IDIC Formulation	5
2.1	Handling the Ill-posedness of the Inverse Problem	6
2.2	Inexact Newton-CG Method for the Regularized Inverse Problem	7
2.3	Derivation of the Adjoints and Hessian actions for ∞ -IDIC	8
2.4	Derivations of the gradient and Hessian action for Regularization Terms	9
2.4.1	Primal-Dual TV	10

*Corresponding author
Email address: kirchhoff@my.utexas.edu (Joseph Kirchhoff)

3	Numerical Results	10
3.1	Introducing the PDEs	10
3.1.1	Linear Elasticity	11
3.1.2	Neo-Hookean Hyperelasticity	11
3.2	Spatially Varying Inversions of Heterogeneous Fields	12
3.3	Regularization Mitigates the Ill-Posedness	13
3.4	Mesh-independent Performance	14
3.5	∞ -IDIC Robustness	16
3.5.1	Image Noise and Force Measurement Error	17
3.5.2	Relationship Between Speckle Size and Resolution of Inferred Features	17
3.5.3	Varying the Forcing Condition	18
3.5.4	An Information Gain Heuristic	19
3.5.5	Information Gain Increases with Refining Speckle	20
3.5.6	Information Gain Increases with Larger Forcing	20
3.6	Useful Extensions of ∞ -IDIC	21
3.6.1	Learning More Through Multiple Experiments	21
3.6.2	Post-processing for Modulus-Induced Stress Concentrations	22
4	Conclusion	23
5	Acknowledgments	24
Appendix A	Details on Numerical Implementation	27
Appendix B	Bayesian ∞-IDIC and expected information gain	28
Appendix B.1	Notes on the prior and the function space limit	30

1. Introduction

We present a novel formulation of inverse problems for spatially-varying parameter fields of continuum materials using images of their deformations. In particular, we simultaneously pose the inversion and the image registration problems as a continuum limit in a function space setting. This formulation allows for the derivation of scalable and efficient *discretization dimension-independent* algorithms for the inversion of heterogeneous material properties such as Lamé parameters while maintaining sharpness of features in their reconstruction, via the use of appropriate regularization methods. Our approach aims to address the imperative to characterize materials with spatially-varying properties. In particular we seek the non-destructive identification of arbitrary damage, defects, degradation, anomalies and inclusion without knowledge of their morphology or strength. The need for such methods is laid out in the vision for Materials Testing 2.0 (MT2.0) [1, 2], which calls for a paradigm shift in material characterization within research laboratories, particularly for materials like tissues, composites, welds, foams, etc., where full-field measurements should be used to identify spatially varying properties [2].

Since heterogeneous material defects are not directly observable, their identification must be posed as an inverse problem for arbitrary infinite-dimensional fields. In order to infer such infinite-dimensional fields, a significant amount of *informative* observational data is required. Digital image correlation (DIC) derives displacement data via statistical correlation in differences in speckle patterned images [3]. However since these statistical algorithms do not obey the governing mechanical equilibrium equations, they may lead to cascading errors in the inversion. Integrated DIC (IDIC) introduced by Besnard, Hild and Roux, overcame this issue by constraining the displacements to themselves follow the governing equations of the deformation process [4]. IDIC in its current formulation however inverts for material properties as homogeneous constants *and not* heterogeneous spatial fields [5].

In this paper, we extend the formulation of IDIC to the function space setting, where we formulate the image registration and heterogeneous material inference problem as one monolithically coupled inverse problem. In this setting, we model the image data, and the material properties as infinite-dimensional spatial fields alongside the state variable (e.g., strain). Since the material properties are now formally infinite-dimensional fields, this leads to significant ill-posedness. In order to overcome this ill-posedness we

consider various regularization models, such as $L^2(\Omega)$, $H^1(\Omega)$ Tikhonov regularization for smooth parameter reconstruction, and (primal-dual) total variation (TV) for reconstruction of sharp features. These modeling choices allow us to derive scalable and efficient Newton methods for the solution of these corresponding inverse problems. We term the methods presented herein Infinite-Dimensional IDIC, abbreviated as ∞ -IDIC.

We illustrate our approach through nondestructive evaluation problems for solids where one has observations of undeformed and deformed states of the material, which are painted speckle patterns.

In the numerical results we demonstrate ∞ -IDIC considering linear elastic and neo-Hookean hyperelastic models. We demonstrate that ∞ -IDIC can recover spatially-varying material properties with remarkable accuracy and sharpness; the corresponding stress fields can be post-processed for. This shows promise for predicting failure mechanisms of stress concentrations that arise from spatially varying material properties. Moreover, that our algorithms exhibit dimension-dependent convergence properties, and robustness to observational and loading noise. Our findings show that the information gained by the inversion depends on the loading conditions, which can be exploited by incorporating multiple experiments in a single ∞ -IDIC inverse problem.

1.1. Related works

IDIC has been demonstrated using a wide range of PDEs such as phase-field fracture modeling [5], elasto-plasticity [6–8], anisotropic linear elasticity [9], to name a few. The current research in IDIC focuses on low-dimensional inversion, where the problem’s ill-posedness may not be as severe as in high-dimensional spaces. However, often Tikhonov regularization is opted for or no regularization at all [7]. Rokos et al. address low-dimensional parameter estimation limitations by using IDIC to infer a homogenized modulus and then employing the principle of virtual work for heterogeneous behavior inference [10]. Such studies underscore the necessity for a direct approach to handle IDIC in high-dimensional function spaces.

Unlike IDIC, which directly solves the inverse problem for material parameters, an alternative approach involves using image registration DIC software to acquire a displacement field, subsequently treating it as observational data. Initially devised to validate Finite Element Model (FEM) predictions against DIC observations [11], a technique known as Finite Element Model Updating (FEMU), has evolved into the most commonly used method for material parameter inference from DIC data [12–17]. It was shown that FEMU and IDIC behave similarly in a simple experiment for inferring elasto-plastic parameters [6]. However, IDIC outperforms FEMU in challenging inverse problems, such as those involving image noise, complex loading conditions or experimental errors [18]. FEMU is useful when the DIC algorithm accurately captures the displacement field, yet any error in DIC image registration can compromise the material inversion [19].

Existing IDIC works infer scalar, homogeneous material parameters, largely neglecting spatially varying fields, specifically heterogeneous materials [1, 2]. There is limited work in FEMU that invert for spatially varying material properties for one or two stiff inclusions in soft bodies for tumor identification (elastography) [20–25]. The existing work is reliant on an initial DIC algorithm to provide accurate displacement fields and do not appear to be mesh-independent. Scaling to high dimensional fields is a challenge we address in this paper for inversion of complex material fields. Infinite dimensional inverse problems, where the input parameters are *inferred* from noisy observations of the data, are inherently ill-posed [26, 27]. We solve this using an inverse problem formulation that is familiar to a vast literature. For more information we refer the reader to other works [26–30] for a more in depth treatment of the mathematical aspects of inverse problems.

1.2. Contributions

We present ∞ -IDIC, an image based parameter estimation inverse problem established in function spaces, thereby enabling inversion of spatially varying material properties. We note that the formulation of ∞ -IDIC is general and applies across various materials like tissues, composites, welds, and foams, improving material characterization through full-field measurements. Inferring detailed parameter and stress fields offers new possibilities, a notable advancement for both FEM validation and non-destructive evaluation.

Our overall approach of simultaneously posing the image registration and heterogeneous material inference problems can be extended to other fields such as in fluids (particle tracking velocimetry) [31–36], rheometry (micro-rheometry) [37], medical (image registration) [38–43], robotics (point-set registration) [44–48], and even more broadly for optical flow [49–51]. Figure 1 illustrates the key contributions of our work, by highlighting the novelty of handling heterogeneous material properties, spatial modulus inference and ultimately an accurate stress inference. The key contributions are summarized:

	physics-based registration	heterogeneity	mesh independence
FEMU (scalar) [6, 12–17]	✗	✗	✗
FEMU (field) [20–25]	✗	✓	✗
IDIC (scalar) [4–9]	✓	✗	✗
∞ -IDIC (ours)	✓	✓	✓

Table 1. Summary of different DIC-based material property inverse methods.

1. Formulating ∞ -IDIC in infinite-dimensional function spaces, enabling a dimension-independent algorithm to infer spatially varying parameter fields.
2. Inverting for high-dimensional modulus fields captures heterogeneous material properties.
3. Post-processing for stress fields from the inferred modulus and displacement fields, enabling the identification of heterogeneity induced stress concentrations.

We handle the inherent ill-posedness of the formulated inverse problem via mesh-independent regularization strategies, including $L^2(\Omega)$, $H^1(\Omega)$, primal-dual total variation that result in modeling benefits (e.g., sharpness, smoothing). This results in a regularized inverse problem which we choose to solve with an inexact Newton-CG method using a Gauss–Newton approximation of the Hessian. We demonstrate that the inferred modulus and displacement fields can be used to predict stress fields, enabling the identification of stress concentration-based failure mechanisms. In our numerical experiments, linear elastic and hyperelastic material models are considered; these results showcase the robustness of ∞ -IDIC to noise, mesh independence, and boundary condition variations for complex material systems. We highlight our contributions relative to existing works in Table 1, and illustrate the capabilities of ∞ -IDIC in Figure 1.

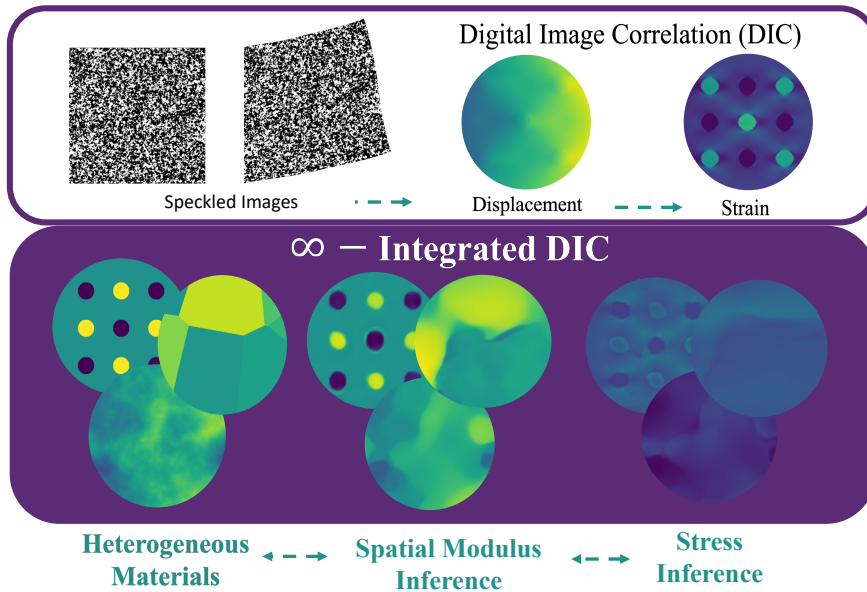


Fig. 1. Our contributions include the development of an infinite-dimensional IDIC framework that enables the inversion of spatially varying parameter fields engineered for heterogeneous materials. Novel contributions are highlighted: handling heterogeneity, infinite dimensional framework, inference for high dimensional modulus fields, and stress fields.

1.3. Layout of the paper

The rest of the paper is divided into two overarching sections: formulation of ∞ -IDIC and numerical results. The formulation section begins with describing the infinite-dimensional inverse problem and its ill-posedness. This is treated using regularization methods, including $L^2(\Omega)$ Tikhonov regularization, $H^1(\Omega)$

Tikhonov smoothing regularization, and total variation to favor sharp reconstructions. We then present an inexact-Newton algorithm for solving the resulting PDE-constrained optimization problems, including the computation of the gradient and (approximation) Hessian actions via adjoint methods. The numerical results’ section showcases the performance of the algorithm in inferring high-dimensional modulus fields. The results are presented for linear elasticity and hyperelasticity, demonstrating the robustness of the algorithm to noise, mesh independence, and boundary condition variations. The paper concludes with a discussion on the results and future work.

2. The ∞ -IDIC Formulation

This section presents the formulation for ∞ -IDIC: the simultaneously coupled infinite-dimensional integrated image registration and parameter identification problem from image data. We assume a given model for the underlying physics (e.g., linear elasticity and hyperelasticity), and assume that the observational data for the inverse problem consist of images of the undeformed and deformed configurations of the body of interest. We assume that the body is painted with a speckle pattern (speckled) to maximize the information about the point wise deformation state via contrast in pixel values (see Figure 2). We then formulate a simultaneously coupled image registration and heterogeneous material property inverse problem, which constrains the admissible displacements to obey the governing conservation laws. This formulation thereby leads to physically consistent reconstruction of both the displacement and heterogeneous parameter fields. We follow the “optimize-then-discretize” (OTD) formalism, and derive all of our proposed methods for the continuum limit of the problems of interest in their corresponding function space setting [52, 53]. This allows us to mitigate the effects of discretization on the performance of the methods in practice.

We consider physical models governing the motion of deforming elastic bodies, formulated as partial differential equations (PDEs). Given a physical domain $\Omega \subset \mathbb{R}^d$ ($d = 2, 3$) and boundary Γ of the (undeformed) solid body, the PDE model defines the mapping from a spatially varying parameter field $m \in \mathcal{M}$ and the traction $t \in \mathcal{T}$ to the state variable $u \in \mathcal{U} := u_0 + \mathcal{V}$ representing the displacement field, where \mathcal{M} , \mathcal{T} , \mathcal{U} are the function spaces for parameter, traction, and displacement fields, respectively. A schematic of an example setup is shown in Figure 2. In particular, \mathcal{M} and \mathcal{T} are assumed to be Hilbert spaces such as the Sobolev spaces $H^k(\Omega)$ and $L^2(\Gamma)$. The state space \mathcal{U} is given as an affine shift of a Hilbert space \mathcal{V} by a finite energy lift of any displacement boundary conditions, u_0 , where u_0 satisfies the desired Dirichlet boundary conditions and $\|u_0\|_{\mathcal{V}} < \infty$. The PDE model can then be written abstractly as

$$\text{PDE Model:} \quad R(u, m, t) = 0, \quad (1)$$

where the PDE residual, $R : \mathcal{U} \times \mathcal{M} \times \mathcal{T} \rightarrow \mathcal{V}'$, is a possibly nonlinear combination of differential operators, and \mathcal{V}' is the (topological) dual of \mathcal{V} . Alternatively, we can formulate the PDE problem in its weak form using \mathcal{V} as the test space,

$$\text{Find } u \in \mathcal{U} \text{ such that} \quad r(u, m, t, v) = \langle R(u, m, t), v \rangle_{\mathcal{V}} = 0, \quad \forall v \in \mathcal{V}, \quad (2)$$

where $r : \mathcal{U} \times \mathcal{M} \times \mathcal{T} \times \mathcal{V} \rightarrow \mathbb{R}$ is the weak residual, and $\langle \cdot, \cdot \rangle_{\mathcal{V}}$ denotes the duality pairing between \mathcal{V}' and \mathcal{V} . Note that $r(u, m, t, v)$ is linear with respect to the test function v . We will assume that the PDE problem is well-posed and admits solution operator $u = u(m, t)$ that is differentiable with respect to the PDE parameters.

In the simplest case, the observational data is given in the form of two images: I_0 and I_1 , capturing the undeformed and deformed specimen, respectively, see Figure 2 for a schematic. This case generalizes to many images in the case of multiple deformation states to better inform the inversion. Mathematically, we consider images as functions of spatial coordinates that return the grayscale pixel value of the image at a particular point, where 0 is black and 255 is white. In particular images I_0 and I_1 are defined over an image domain Ω_I , which is large enough to capture both the undeformed and deformed specimen. Moreover, we consider the space of images to be $\mathcal{I} = H^1(\Omega_I)$ such that they admit at least one spatial derivative.

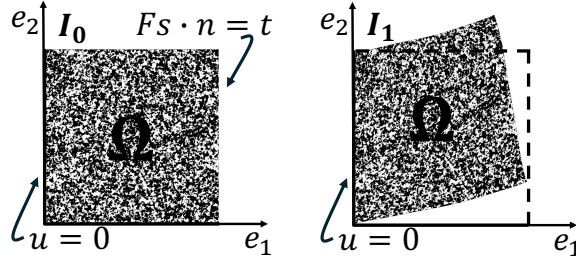


Fig. 2. Schematic of the numerical setup for generating the reference and deformed states. The domain, Ω , is speckled. A traction condition is applied to the right boundary, Γ_R and a clamped condition is applied to the left boundary, Γ_L .

The objective of the coupled inversion and registration is then to find the parameter field m such that the predicted displacement yields an image that matches the image of the deformed specimen. This is formulated using the following minimization problem,

$$\min_{m \in \mathcal{M}, u \in \mathcal{U}} \mathcal{J}(u, m; I_0, I_1) := \Phi(u; I_0, I_1) + \mathcal{R}(m) \text{ subject to } R(u, m, t) = 0,$$

Here, the image misfit

$$\Phi(u; I_0, I_1) = \frac{1}{2} \int_{\Omega} |I_1(x + u(x)) - I_0(x)|^2 dx, \quad (3)$$

compares the pixel values of the images, where $|\cdot|$ denotes the magnitude. Simply put, the misfit is the difference in the two images in function space where the displaced image is ‘pulled-back’ by the displacement field.

For the elliptic and parabolic PDEs that govern the deformation of solid bodies, the mapping from material parameter fields m to the displacement u tends to be smoothing, i.e., highly oscillatory perturbations of m do not significantly affect u . Thus, the inversion of full heterogeneous parameter fields is inherently ill-posed. To overcome this ill-posedness, we additionally introduce regularization on the parameter field, $\mathcal{R}(m)$, to the cost functional \mathcal{J} .

It is important to note that, unlike in the case of DIC, where the ill-posedness is in directly identifying the displacement field based on pixel correlations, the displacement field in our formulation is completely constrained by the underlying PDE solution, $u = u(m, t)$. Instead, the regularization is prescribed on the parameter field to overcome the ill-posedness arising from the desire to invert for heterogeneous material parameters.

2.1. Handling the Ill-posedness of the Inverse Problem

As previously discussed, it is typical for the PDE solution operator to be insensitive to certain directions (subspaces) in the parameter space. In such settings, observation data alone is not sufficient to identify the material parameters in these subspaces, and regularization is required to make the inverse problem well-posed. The choice of regularization \mathcal{R} is based on the prior knowledge about the parameter field, and is tightly related to the choice of a prior distribution on the parameters in Bayesian inference. Thus, the context should motivate the decision of which regularization to use.

In this work, we consider three choices of regularization terms. In the simplest case, we can consider L^2 Tikhonov regularization to penalize the deviation of m from a reference value $\bar{m} \in \mathcal{M}$,

$$\mathcal{R}_{L^2}(m) = \frac{\gamma_{L^2}}{2} \int_{\Omega} |m(x) - \bar{m}(x)|^2 dx. \quad (4)$$

where $\gamma_{L^2} > 0$ is a weighting factor. For example, one can take $\bar{m}(x) \equiv m_{\text{nominal}}$, where m_{nominal} is a nominal value for the material under consideration, maybe obtained from a material database. In addition to directly regularizing the values of m , another option is to penalize steeper gradients of m to impose a preference for smooth parameter fields. To this end, H^1 Tikhonov regularization uses an $L^2(\Omega)$ penalization on the gradient,

$$\mathcal{R}_{H^1}(m) = \frac{\gamma_{H^1}}{2} \int_{\Omega} \|\nabla m(x)\|_2^2 dx, \quad (5)$$

with weighting $\gamma_{H^1} > 0$, where $\|\cdot\|_2$ denotes the Euclidean (ℓ^2) norm of a vector in \mathbb{R}^d . The smoothing effect introduced by H^1 regularization tends to blur sharp edges.

If it is expected *a priori* that there are sharp edges (i.e., voids, fibers, particles) then total variation (TV) regularization can be introduced. Unlike H^1 , TV regularization is defined using a $L^1(\Omega)$ penalization on the gradient

$$\mathcal{R}_{TV}(m) = \gamma_{TV} \int_{\Omega} \|\nabla m(x)\|_2 dx, \quad (6)$$

with weighting $\gamma_{TV} > 0$. Since the $L^1(\Omega)$ norm is sparsifying, TV regularization tends to result in piecewise constant parameter fields. However, $\mathcal{R}_{TV}(m)$ is not differentiable when $\nabla m = 0$. A common approach to addressing this issue is by introducing a smoothing parameter, $\epsilon > 0$, to define a smoothed TV regularization term,

$$\mathcal{R}_{TV}^{\epsilon}(m) = \gamma_{TV}^{\epsilon} \int_{\Omega} \sqrt{\|\nabla m(x)\|_2^2 + \epsilon} dx, \quad (7)$$

with weighting $\gamma_{TV}^{\epsilon} > 0$. This smooths the TV functional at zero, making it differentiable and also leads to a more positive definite Hessian of $\mathcal{R}_{TV}^{\epsilon}$. However, this will tend to produce smooth transitions between piecewise constant regions, where the sizes of the transition zones increase with ϵ . Reducing ϵ brings the approximation $\mathcal{R}_{TV}^{\epsilon}$ closer to $\mathcal{R}_{TV}(m)$, effectively sharpening edges, but also makes the Hessian less positive-definite.

For the Tikhonov regularization terms, one tunes only the weighting factor (γ_{L^2} or γ_{H^1}) as a hyperparameter, whereas in the case of smoothed TV regularization, there are two hyperparameters, γ_{TV} and ϵ . In this study, ϵ is systematically reduced to sharpen the edges of the parameter field, until the problem becomes ill-conditioned again, and the optimization algorithm fails to converge. Typically, one will use a combination of regularization terms, such as combining L^2 with H^1 or TV, leading to two or three tunable hyperparameters. A common approach is to use a heuristic, such as the L-curve criterion, to determine the optimal values of the regularization parameters. Regardless of method, it is important to tune the regularization to result in a well-posed optimization problem that considers the physical understanding of the problem at hand.

2.2. Inexact Newton-CG Method for the Regularized Inverse Problem

The incorporation of the regularization terms into the inverse problem results in a well-posed optimization problem in which we solve with a scalable and efficient inexact Newton-CG method. To this end, we opt for a reduced-space approach, in which the minimization problem (3) is solved by explicitly eliminating the PDE constraint $R(u, m, t) = 0$ using the solution operator $u = u(m, t)$. This leads to an unconstrained minimization problem,

$$\min_{m \in \mathcal{M}} \widehat{\mathcal{J}}(m) := \mathcal{J}(u(m, t); I_0, I_1) = \Phi(u(m, t); I_0, I_1) + \mathcal{R}(m), \quad (8)$$

since u directly depends on m through the solution of the PDE. To solve (8), we opt for a gradient based optimization scheme. Notably, gradient descent is not mesh-independent since it is not affine-invariant. Instead, we consider Newton's method for solving (8) since its affine-invariance leads to dimension independent performance [54]. Starting at an initial guess, m^0 , m^k is updated at each step for $k = 1, 2, \dots$, by

$$m^{(k+1)} = m^{(k)} + \beta^{(k)} \delta m^{(k)}. \quad (9)$$

where $\beta^{(k)}$ is the step size. The search direction, $\delta m^{(k)}$, is determined through the Newton step, written below in variational form,

$$\left\langle D_{mm} \widehat{\mathcal{J}}(m^{(k)}) \delta m^{(k)}, \tilde{m} \right\rangle_{\mathcal{M}} = - \left\langle D_m \widehat{\mathcal{J}}(m^{(k)}), \tilde{m} \right\rangle_{\mathcal{M}} \quad \forall \tilde{m} \in \mathcal{M}, \quad (10)$$

where $D_m \widehat{\mathcal{J}}$ and $D_{mm} \widehat{\mathcal{J}}$ are the first and second Gâteaux derivatives (i.e., gradient and Hessian respectively) of the (reduced-space) cost functional. These derivatives, which are implicitly defined through the PDE solution operator, can be computed efficiently using the adjoint method; we derive their explicit forms in section 2.3. However, explicitly constructing the full Hessian is intractable since a pair of linearized forward and adjoint PDEs must be solved to form each column of the Hessian via Hessian-vector products. The number of columns scales with the discretization dimension, not the *inherent* dimension of the problem.

Instead, we leverage the inexact Newton-CG (INCG) algorithm [55, 56], which is a matrix-free Krylov method for the approximation solution of the Newton system (10), only requiring Hessian actions, thereby avoiding explicitly forming and inverting the Hessian. Using the gradient norm for the inexact solves,

$$\left\| D_m \widehat{\mathcal{J}}(m^{(k)}) + D_{mm} \widehat{\mathcal{J}}(m^{(k)}) \delta m^{(k)} \right\|_{\mathcal{M}'} \leq r_{tol} \left\| D_m \widehat{\mathcal{J}}(m^{(k)}) \right\|_{\mathcal{M}'} \quad (11)$$

allows one to avoid expensive computations far away from an optimum where only descent is needed. As the iterates proceed the tolerance grows tighter and better Newton iterates are produced. We use the tolerance, $r_{tol} = \min \left(0.5, \sqrt{\widehat{\mathcal{J}}(m^{(k)}) / \widehat{\mathcal{J}}(m^{(0)})} \right)$, as this theoretically allows for superlinear convergence in local regimes [57]. Thus, with the chosen CG termination criterion, inexact Newton-CG ensures that only few Hessian actions are used to solve for the search direction in the global (pre-asymptotic) regime, where the descent does not benefit significantly from Hessian information. On the other hand, in the regime of local convergence, (10) is solved to a finer tolerance, incorporating more accurate curvature information which classically leads to superlinear convergence when the true Hessian is used. In this work we utilize the Gauss–Newton Hessian instead due to our assumed smoothness of the image field, so we settle for fast linear convergence.

2.3. Derivation of the Adjoint and Hessian actions for ∞ -IDIC

In this section, we derive the gradient and Hessian actions of the (reduced-space) cost functional that are required for the INCG algorithm. To compute the gradient of \mathcal{J} , we employ a Lagrangian approach by introducing an adjoint variable, $p \in \mathcal{V}$ and a Lagrangian functional,

$$\mathcal{L}(u, m, t, p) = \Phi(u; I_0, I_1) + \mathcal{R}(m) + r(u, m, t, p). \quad (12)$$

The m -reduced gradient is formed by eliminating the Karush–Kuhn–Tucker (KKT) conditions associated with p and m , i.e., we set the first variations of the Lagrangian with respect to p and u to zero, such that the first derivative of the cost functional with respect to m , $D_m \mathcal{J}$, coincides with $\partial_m \mathcal{L}$. For brevity, we will omit the arguments to \mathcal{L} , where it is understood that they are to be evaluated at (u, m, t) . The forward equation, adjoint equation, and gradient equation are thus given by,

$$r(u, m, t, \tilde{p}) = 0 \quad \forall \tilde{p} \in \mathcal{V}, \quad (13a)$$

$$\langle \partial_u r(u, m, t, p), \tilde{u} \rangle_{\mathcal{V}} = -\langle \partial_u \Phi, \tilde{u} \rangle_{\mathcal{V}} \quad \forall \tilde{u} \in \mathcal{V}, \quad (13b)$$

$$\langle D_m \mathcal{J}, \tilde{m} \rangle_{\mathcal{M}} = \langle \partial_m \mathcal{R}(m), \tilde{m} \rangle_{\mathcal{M}} + \langle \partial_m r(u, m, t, p), \tilde{m} \rangle_{\mathcal{M}} \quad \forall \tilde{m} \in \mathcal{M}, \quad (13c)$$

respectively, where we can write out the misfit’s contribution to the adjoint equation as

$$\langle \partial_u \Phi, \tilde{u} \rangle_{\mathcal{V}} = \int_{\Omega} (I_1(x + u(x)) - I_0(x)) \nabla I_1(x + u(x)) \cdot \tilde{u}(x) dx \quad \forall \tilde{u} \in \mathcal{V}. \quad (14)$$

The gradient of the cost function can therefore be found by the subsequent process of first solving (13a) for u , then given u solving (13b) for p , then given u and p forming the variational gradient in (13c). To derive the Hessian action, we introduce the Hessian meta-Lagrangian,

$$\mathcal{L}^H(u, m, t, p, \hat{u}, \hat{m}, \hat{p}) = \langle \partial_u \mathcal{L}, \hat{u} \rangle_{\mathcal{V}} + \langle \partial_m \mathcal{L}, \hat{m} \rangle_{\mathcal{M}} + \langle \partial_p \mathcal{L}, \hat{p} \rangle_{\mathcal{V}}. \quad (15)$$

Analogous to the gradient formation procedure, we set $\partial_p \mathcal{L}^H = 0$, $\partial_u \mathcal{L}^H = 0$, such that $\langle D_{mm} \mathcal{J} \hat{m}, \tilde{m} \rangle_{\mathcal{M}} = \langle \partial_m \mathcal{L}^H, \tilde{m} \rangle_{\mathcal{M}}$, where $D_{mm} \mathcal{J} \hat{m}$ is the Hessian acting in a direction $\hat{m} \in \mathcal{M}$. These give rise to the incremental forward, incremental adjoint, and the Hessian action equations,

$$\langle \partial_u r(u, m, t, \tilde{p}), \hat{u} \rangle_{\mathcal{V}} = -\langle \partial_m r(u, m, t, \tilde{p}), \hat{m} \rangle_{\mathcal{M}} \quad \forall \tilde{p} \in \mathcal{V}, \quad (16a)$$

$$\langle \partial_u r(u, m, t, \tilde{p}), \tilde{u} \rangle_{\mathcal{V}} = -(\langle \partial_{uu} \Phi \hat{u}, \tilde{u} \rangle_{\mathcal{V}} + \langle \partial_{mu} r(u, m, t, p) \hat{m}, \tilde{u} \rangle_{\mathcal{V}} + \langle \partial_{uu} r(u, m, t, p) \hat{u}, \tilde{u} \rangle_{\mathcal{V}}) \quad \forall \tilde{u} \in \mathcal{V}, \quad (16b)$$

$$\begin{aligned} \langle D_{mm} \mathcal{J} \hat{m}, \tilde{m} \rangle_{\mathcal{M}} &= \langle \partial_{mm} \mathcal{R}(m) \hat{m}, \tilde{m} \rangle_{\mathcal{M}} + \langle \partial_{mm} r(u, m, t, \tilde{p}) \hat{m}, \tilde{m} \rangle_{\mathcal{M}} \\ &\quad + \langle \partial_m r(u, m, t, \tilde{p}) \hat{m}, \tilde{p} \rangle_{\mathcal{M}} + \langle \partial_{um} r(u, m, t, p) \hat{u}, \tilde{m} \rangle_{\mathcal{M}} \end{aligned} \quad \forall \tilde{m} \in \mathcal{M}, \quad (16c)$$

respectively.

The incremental forward equation (16a) is solved for \hat{p} , the incremental adjoint problem (16b) is solved for \hat{u} , and the variables are combined with u and p to give the Hessian of the cost functional acting in the direction of \hat{m} using (16c). Again, we can explicitly write out the misfit contribution to the incremental adjoint equation as

$$\begin{aligned} \langle \partial_{uu} \Phi \hat{u}, \tilde{u} \rangle_{\mathcal{V}} &= \int_{\Omega_0} (I_1(x + u(x)) - I_0(x)) \nabla^2 I_1(x + u(x)) \hat{u} \cdot \tilde{u} \, dx \\ &\quad + \int_{\Omega_0} (\nabla I_1(x + u(x)) \cdot \hat{u}) (\nabla I_1(x + u(x)) \cdot \tilde{u}) \, dx \quad \forall \hat{u}, \tilde{u} \in \mathcal{V}. \end{aligned}$$

The second spatial derivative of an image, specifically $\nabla^2 I_1(x + u(x))$, may not exist given our assumed \mathcal{C}^1 smoothness. To avoid more restrictive assumptions of \mathcal{C}^2 smoothness, we employ a Gauss–Newton approximation of the Hessian [54], in which we drop the $(I_0(x) - I_1(x + u(x))) \nabla^2 I_1(x + u(x)) \hat{u} \cdot \tilde{u}$ term, leaving only the terms involving ∇I_1 . That is, we take the approximation

$$\langle \partial_{uu} \Phi \hat{u}, \tilde{u} \rangle_{\mathcal{V}} \approx \int_{\Omega_0} (\nabla I_1(x + u(x)) \cdot \hat{u}) (\nabla I_1(x + u(x)) \cdot \tilde{u}) \, dx \quad \forall \hat{u}, \tilde{u} \in \mathcal{V}. \quad (17)$$

This approximation is accurate when the data misfit $I_1(x + u(x)) - I_0(x)$ is small, but since $I_1 \not\approx I_0$ due to the observational noise, the Gauss–Newton Hessian does not converge to the true Hessian. We do not expect the typical superlinear convergence of inexact Newton methods and instead fast linear convergence.

Remark 2.1. It is possible that the true Hessian can be computed if the second spatial derivative of the image is well-defined, but this is not the case for the images considered in this work as they are piece-wise constant speckle patterns where the second derivative is undefined at the edges of the speckles.

2.4. Derivations of the gradient and Hessian action for Regularization Terms

In this section we derive the gradient and Hessian actions of the regularization terms. Beginning with the L^2 Tikhonov regularization, we have the following gradient and Hessian action:

$$\langle D_m \mathcal{R}_{L^2}(m), \tilde{m} \rangle_{\mathcal{M}} = \gamma_{L^2} \int_{\Omega} m \tilde{m} \, dx \quad \forall \tilde{m} \in \mathcal{M}, \quad (18a)$$

$$\langle D_{mm}^2 \mathcal{R}_{L^2}(m) \hat{m}, \tilde{m} \rangle_{\mathcal{M}} = \gamma_{L^2} \int_{\Omega} \hat{m} \tilde{m} \, dx \quad \forall \tilde{m}, \hat{m} \in \mathcal{M}. \quad (18b)$$

For the H^1 regularization, we have the following gradient and Hessian action:

$$\langle D_m \mathcal{R}_{H^1}(m), \tilde{m} \rangle_{\mathcal{M}} = \gamma_{H^1} \int_{\Omega} \nabla m \cdot \nabla \tilde{m} \, dx \quad \forall \tilde{m} \in \mathcal{M}, \quad (19a)$$

$$\langle D_{mm}^2 \mathcal{R}_{H^1}(m) \hat{m}, \tilde{m} \rangle_{\mathcal{M}} = \gamma_{H^1} \int_{\Omega} \nabla \hat{m} \cdot \nabla \tilde{m} \, dx \quad \forall \tilde{m}, \hat{m} \in \mathcal{M}. \quad (19b)$$

For the TV regularization we have the following gradient and Hessian action, which we refer to as the primal formulation of TV:

$$\langle D_m \mathcal{R}_{TV}^{\epsilon}(m), \tilde{m} \rangle_{\mathcal{M}} = \gamma_{TV}^{\epsilon} \int_{\Omega} \frac{\nabla m}{\sqrt{\|\nabla m\|_2^2 + \epsilon}} \cdot \nabla \tilde{m} \, dx \quad \forall \tilde{m} \in \mathcal{M}. \quad (20a)$$

$$\langle D_{mm}^2 \mathcal{R}_{TV}^{\epsilon}(m) \hat{m}, \tilde{m} \rangle_{\mathcal{M}} = \gamma_{TV}^{\epsilon} \int_{\Omega} \frac{1}{\sqrt{\|\nabla m\|_2^2 + \epsilon}} \left[\left(I - \frac{\nabla m \otimes \nabla m}{\|\nabla m\|_2^2 + \epsilon} \right) \nabla \hat{m} \right] \cdot \tilde{m} \, dx \quad \forall \tilde{m}, \hat{m} \in \mathcal{M}, \quad (20b)$$

where $\|\cdot\|_2$ denotes the usual Euclidean norm for vectors in \mathbb{R}^d . It has been shown that anisotropy in the $(I - \frac{\nabla m \otimes \nabla m}{\|\nabla m\|_2^2 + \epsilon})$ term causes convergence issues; as ϵ is reduced the radius of local convergence also is reduced and the method may not converge globally [58]. Additionally, the ϵ parameter is mesh dependent, so one must retune the hyperparameter as the mesh is refined. A first attempt to mitigate this is to scale ϵ with the mesh size [59]. However, this method proved unsuccessful in our experiments and mesh-independence was not achieved, possibly due to the nonlinear nature of the ∞ -IDIC formulation (and nonlinear PDE used in demonstration). To mitigate the mesh-dependence, we instead employ a primal-dual approach which empirically leads to a mesh-independent algorithm.

2.4.1. Primal-Dual TV

The issues with convergence of Newton's method for the primal $\mathcal{R}_{TV}^\epsilon$ are due to the highly ill-conditioned Hessian term (20b). The ill-conditioning of the Hessian is due to taking an additional derivative of the $\frac{\nabla m}{\sqrt{\|\nabla m\|_m^2 + \epsilon}}$ term in the gradient (20a). In [58], the authors proposed a solution to this problem via an equivalent *primal-dual* formulation of the TV-regularized optimization problem. In this formulation an additional dual variable $w = \frac{\nabla m}{\sqrt{\|\nabla m\|_m^2 + \epsilon}}$ is introduced leading to the primal-dual interpretation of TV:

$$\mathcal{R}_{TV}^\epsilon(m) = \gamma_{TV} \sup_{w \in \mathcal{W}} \left\{ \int_{\Omega} m(x) \nabla \cdot w(x) dx : \|w(x)\|_2 \leq 1 \quad \forall x \in \Omega \right\}, \quad (21)$$

where \mathcal{W} is an appropriately chosen function space for the dual variable. For example when approximating \mathcal{M} by continuous Galerkin finite elements with order r , $\mathcal{CG}^r(\Omega)$ the space of $r - 1^{th}$ order discontinuous Galerkin finite elements, $\mathcal{DG}_r(\Omega)^d$, is a suitable choice for \mathcal{W} , [60, 61]. In the primal-dual TV, the first order optimality condition (13c) for m ,

$$\langle D_m \mathcal{J}, \tilde{m} \rangle_{\mathcal{M}} = 0 \quad \forall \tilde{m} \in \mathcal{M} \quad (22)$$

can be rewritten as a system for the primal-dual pair (m, w) :

$$\gamma_{TV}^\epsilon \int_{\Omega} w \cdot \nabla \tilde{m} dx + \langle \partial_m r(u, m, t, p), \tilde{m} \rangle_{\mathcal{M}} = 0 \quad \forall \tilde{m} \in \mathcal{M}, \quad (23a)$$

$$\int_{\Omega} \left(\sqrt{\|\nabla m\|_2^2 + \epsilon} \right) w \tilde{w} dx - \int_{\Omega} \nabla m \tilde{w} dx = 0 \quad \forall \tilde{w} \in \mathcal{W}, \quad (23b)$$

The subsequent linearization to a Newton system for the pair (m, w) leads to a better-conditioned optimization problem and empirically superior convergence properties [58]. At every Newton iteration, we compute updates for both m and w , as follows:

$$\begin{aligned} m^{(k+1)} &= m^{(k)} + \alpha_m^{(k)} \delta m^{(k)}, \\ w^{(k+1)} &= w^{(k)} + \alpha_w^{(k)} \delta w^{(k)}. \end{aligned}$$

We first solve for $\delta m^{(k)}$ via the Newton system (10), where the Hessian vector product is defined as in (16c), with

$$\langle \partial_{mm} \mathcal{R}_{TV}^\epsilon(m^{(k)}) \delta m^{(k)}, \tilde{m} \rangle_{\mathcal{M}} = \gamma_{TV}^\epsilon \int_{\Omega} \frac{1}{\sqrt{\|\nabla m^{(k)}\|_2^2 + \epsilon}} \left[\left(I - A(m^{(k)}, w^{(k)}) \right) \nabla \tilde{m} \right] \cdot \nabla \delta m^{(k)} dx, \quad (24a)$$

$$A(m, w) = \frac{1}{2} w \otimes \frac{\nabla m}{\sqrt{\|\nabla m\|_2^2 + \epsilon}} + \frac{1}{2} \frac{\nabla m}{\sqrt{\|\nabla m\|_2^2 + \epsilon}} \otimes w. \quad (24b)$$

Once $\delta m^{(k)}$ is computed, $m^{(k+1)}$ is found via a line search to determine an appropriate choice of $\alpha_m^{(k)}$. Next $\delta w^{(k)}$ is found from the following relation:

$$\langle \delta w^{(k)}, \tilde{w} \rangle_{\mathcal{W}} = \int_{\Omega} \frac{1}{\sqrt{\|\nabla m^{(k+1)}\|_2^2 + \epsilon}} \left(I - A(m^{(k+1)}, w^{(k)}) \right) \nabla \tilde{w} - w^{(k)} \tilde{w} + \frac{\nabla m^{(k+1)}}{\sqrt{\|\nabla m^{(k+1)}\|_2^2 + \epsilon}} \tilde{w} dx \quad \forall \tilde{w} \in \mathcal{W}. \quad (25)$$

The dual variable $w^{(k+1)}$ is then updated via a line search for $\alpha_w^{(k)}$, which is chosen such that the condition $\|w(x)\|_2 \leq 1$ is met for all $x \in \Omega$, as in (21). Similar to [58] we observe significantly improved performance over the primal version of TV. In particular, our choice of ϵ is mesh-independent, and we observe consistent and improved convergence across all meshes. With the primal version of TV we were only able to observe convergence on small meshes, after significant hyperparameter tuning to find an appropriate choice of ϵ .

3. Numerical Results

3.1. Introducing the PDEs

We evaluate ∞ -IDIC's efficacy via synthetic data from two PDE models for elastic deformation problems: linear elasticity and hyperelasticity. The parameter fields are complex, spatially varying examples which represent materials with voids, stiff particles, and grain structures. We present the PDEs in the following section, but leave details of the data generation (i.e., image creation and noise corruption) to Appendix A.

3.1.1. Linear Elasticity

We consider unit square physical domains $\Omega = (0, 1)^2$, with spatial coordinates $x \in \Omega$. The strong form of the linear elasticity PDE over Ω is given by

$$\begin{aligned} \nabla \cdot (\sigma(u)) &= 0 & \text{in } \Omega, \\ u &= 0 & \text{on } \Gamma_L, \\ \sigma(u) \cdot n &= t & \text{on } \Gamma_R, \\ \sigma(u) \cdot n &= 0 & \text{on } \Gamma_T \cup \Gamma_B. \end{aligned} \quad (26)$$

We prescribe a fixed displacement condition on the left boundary Γ_L , a traction t on the right boundary Γ_R , and traction-free boundary conditions on top and bottom boundaries, Γ_T and Γ_B (see Figure 2). Here, the stress tensor is given by $\sigma(u) = \lambda \nabla \cdot u I + 2\mu \varepsilon(u)$, where $\varepsilon(u) = \frac{1}{2}(\nabla u + \nabla u^T)$, and λ and μ are the Lamé parameters. We write the traction in terms of its normal and shear components,

$$t = t_{\text{normal}} e_1 + t_{\text{shear}} e_2. \quad (27)$$

The Lamé parameters are related to the Young's modulus, E , and Poisson's ratio, ν , by

$$\lambda = \frac{E\nu}{(1+\nu)(1-2\nu)}, \quad \mu = \frac{E}{2(1+\nu)}. \quad (28)$$

We fix the Poisson's ratio as $\nu = 0.35$ and invert for the Young's modulus, $E(x)$. To prevent the modulus from becoming negative, we parametrize it using the exponential function, i.e., $E(x) = e^{m(x)}$, and invert for the log-Young's modulus m . Using the notation $H_L^1(\Omega) = \{u \in H^1(\Omega) : u|_{\Gamma_L} = 0\}$, the variational form of the PDE (2) is: find $u \in H_L^1(\Omega)$ such that,

$$r_{\text{LE}}(u, m, t, v) := \int_{\Omega} \sigma(u, m) : \varepsilon(v) dx - \int_{\Gamma_R} t \cdot v ds = 0 \quad \forall v \in H_L^1(\Omega). \quad (29)$$

3.1.2. Neo-Hookean Hyperelasticity

As a second PDE problem we consider a hyperelastic material model, which is commonly used to model rubber-like materials, biological tissues, and other materials that undergo large deformations. This problem setting is of interest for elastography [23, 24]. The stress-strain relationship is nonlinear in nature, making it a more difficult numerical test case. Here we use a material point description, unlike linear elasticity where we use spatial points. Let $X \in \Omega$ be the material point where X is related to the spatial point, x , through the displacement, $u(X) = x - X$. In hyperelasticity, the internal forces that develop in the material are derived from a strain energy function, $W = W(X, \mathbf{C}(X))$, where $\mathbf{C}(X) = \mathbf{F}(X)^T \mathbf{F}(X)$ is the right Cauchy–Green deformation tensor and $\mathbf{F}(X) = I + \nabla u(X)$ is the deformation gradient. For the neo-Hookean model, the strain energy function is given by

$$W(X, \mathbf{C}(X)) = \frac{\mu(X)}{2} (\text{tr}(\mathbf{C}(X)) - 3) - \mu(X) \ln J(X) + \frac{\lambda(X)}{2} \ln^2 J(X), \quad (30)$$

where μ and λ are the Lamé parameters, $\text{tr}(\mathbf{C})$ is the trace of \mathbf{C} , and $J = \det(\mathbf{F})$ is the determinant of the deformation gradient. The strong form of hyperelasticity PDE is then given by,

$$\begin{aligned} \nabla \cdot (\mathbf{F}(u) \mathbf{S}(u)) &= 0 & \text{in } \Omega, \\ u &= 0 & \text{on } \Gamma_L, \\ \mathbf{F}(u) \mathbf{S}(u) \cdot n &= t & \text{on } \Gamma_R, \\ \mathbf{F}(u) \mathbf{S}(u) \cdot n &= 0 & \text{on } \Gamma_T \cup \Gamma_B. \end{aligned} \quad (31)$$

where $\mathbf{S}(X) = 2 \frac{\partial W}{\partial \mathbf{C}(X)}$ is the second Piola–Kirchhoff stress tensor, and the applied traction t is the same as in Equation (27).

Similar to linear elasticity, the Young's modulus defined using the parametrization $E(x) = e^{m(x)}$ to ensure positivity and fix the Poisson's ratio as $\nu = 0.35$. Thus, the variational form of the hyperelasticity PDE is: find $u \in H_L^1(\Omega)$ such that for all $v \in H_L^1(\Omega)$,

$$r_{\text{HE}}(u, m, t, v) = \int_{\Omega} \mathbf{S}(u, m) : \varepsilon(v) dx - \int_{\Gamma_R} t \cdot v ds = 0. \quad (32)$$

We use the finite element method to solve both PDE problems. In particular, we employ piecewise linear Lagrangian triangular elements to form finite element spaces $\mathcal{M}^h \approx \mathcal{M}$ and $\mathcal{V}^h \approx \mathcal{V}$. For primal-dual total variation, We discretize the dual variable, w , using 0th order discontinuous Galerkin to form a finite element space $\mathcal{W}^h \approx \mathcal{W}$. We utilize an $N \times N$ mesh, where N is the number of elements per side. We handle all finite element calculations in FEniCS [62] and use hIPPYlib to assist with the adjoint computations [63].

3.2. Spatially Varying Inversions of Heterogeneous Fields

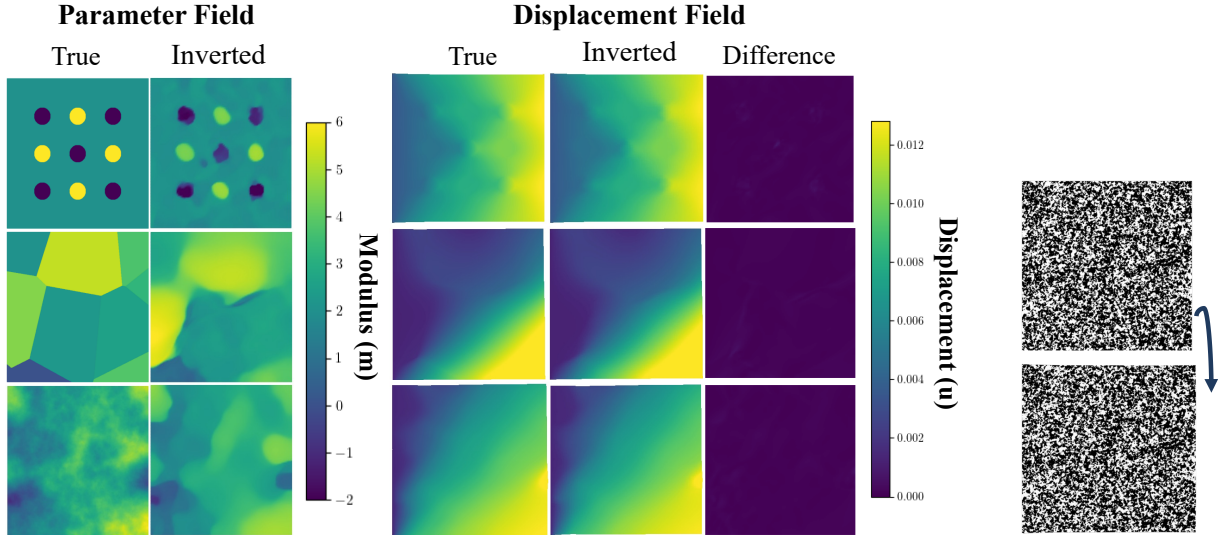


Fig. 3. Inverse results for Linear Elasticity for varying features. Each inversion is based on a simple tensile experiment with $\sim 0.2\%$ strain. 10% and 5% noise were applied to the image brightness and force measurements, respectively. The inverse problem is solved using the primal-dual total variation formulation ($\gamma_{L^2} = 5 \times 10^{-8}$, $\gamma_{TV} = 1 \times 10^{-6}$) with a mesh size of 100×100 , and the speckle correlation length is 0.01. The initial guess was uniformly $m(x) = 2$. The synthetic images are shown on the right.

We begin our numerical experiments by presenting three examples of linear elasticity generated with small strains ($\sim 0.5\%$) in Figure 3. Even at small strains, the inversions capture the spatial distribution well. Notably, the inversions tend to prefer piece wise constant features due to the choice of TV regularization, which is particularly useful when sharp features (i.e., fibers, grains, etc.) are present, as seen in the first two examples. Additionally, we present inversions for hyperelasticity, which could pose greater difficulty due to higher deformations ($\sim 3\%$ strain). In Figure 4, we show five inverse problems ranging from small voids and grain boundaries to cracks. It is suspected that the increase in strain led to improved inversions. Notably, we successfully find features of varying geometries and sizes, as demonstrated by complex blobs, Voronoi tessellations, and a thin, branching crack. In the first case, we detected tiny void regions ($m(x) = -2$) with a radius 1/50th of the domain length; Existing FEMU studies showed inversions for two features with a radius 1/5th of the domain length [23, 25]. However, in all of our inversions, the modulus values do not perfectly match, due to the inverse problem's ill-posedness.

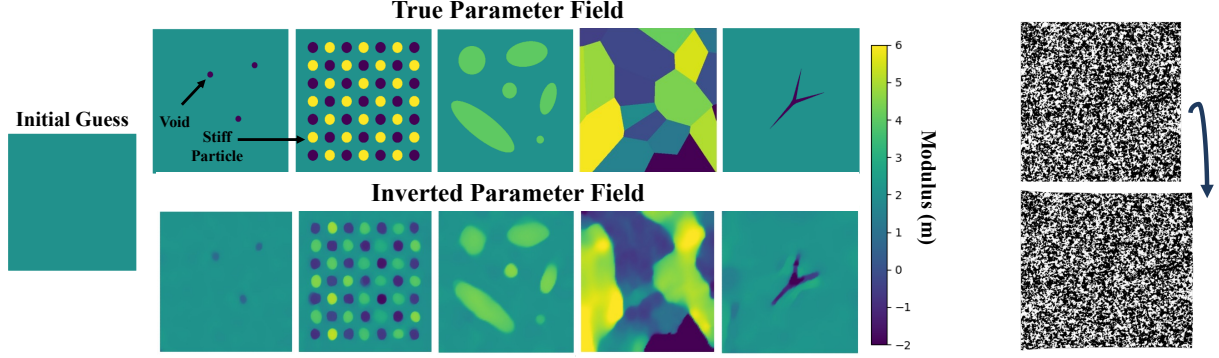


Fig. 4. Inverse results for Hyperelasticity for varying features. Each inversion is based on a simple tensile experiment with $\sim 5\%$ strain. 10% and 5% noise were applied to the image brightness and force measurements, respectively. The inverse problem is solved using the primal-dual total variation formulation ($\gamma_{L^2} = 5 \times 10^{-6}$, $\gamma_{TV} = 7.5 \times 10^{-4}$) with a mesh size of 100×100 , and the speckle correlation length is 0.01.

Ill-posedness likely arises from multiple sources including the non-uniqueness of modulus fields that produce displacement fields satisfying the image data misfit. As seen in Figure 3, the inverted displacement fields almost replicate the true displacement fields despite the modulus field not perfectly matching. There is a non-unique relationship between m and u that satisfies the data, resulting in a non-unique solution to the inverse problem. This effect may also be influenced by the initial guess ($m(x) = 2$), which is two orders of magnitude different from the true solution’s features. The synthetic image and force data were corrupted with noise (see Appendix A) and noise significantly impacts the inversions due to the problem’s ill-conditioning. A key aspect of the ill-posedness is the inherent dissipative nature of elasticity due to the damping of high frequency modes of m by the forward physics [27, 28]. Unaddressed, these high frequency modes would potentially lead to unstable growth which further necessitates the need for regularization. While we incorporate regularization to reduce the ill-posedness, this ultimately introduces a modeling bias into the inversion. For example, in the case of Gaussian random field (Figure 3) with smooth features, the TV regularization tends to sharpen edges.

3.3. Regularization Mitigates the Ill-Posedness

In this section, we study how the choice of regularization affects the inferred parameter fields. In particular, we solve the inverse problem using L^2 , H^1 and TV regularizations for three distinct m_{true} scenarios: a void inclusion, a bump function, and a Gaussian random field. The inverted parameter fields are shown in Figure 5. These results demonstrate that L^2 regularization is insufficient and introduces artifacts, despite considerable tuning efforts for the weighting parameter, yielding noisy and inaccurate inversions incapable of identifying material features. Although L^2 penalizes deviations from a nominal value, and in principle addresses the ill-posedness of the inverse problem, it does not adequately capture the structure of the material fields. In contrast, both H^1 and TV regularization successfully identify features, with varying degrees of smoothing. TV regularization excels in preserving sharp interfaces, as seen in the void inclusion scenario, while H^1 regularization produces smoother inversions without favoring sharp interfaces. Notably, H^1 regularization is better suited than TV for the Gaussian random field, highlighting the importance of adapting regularization based on the nature of the problem. These observations underscore the significance of selecting an appropriate regularization since the choice determines the preferred modulus field characteristics when the noisy data alone is insufficient to recover the true modulus.

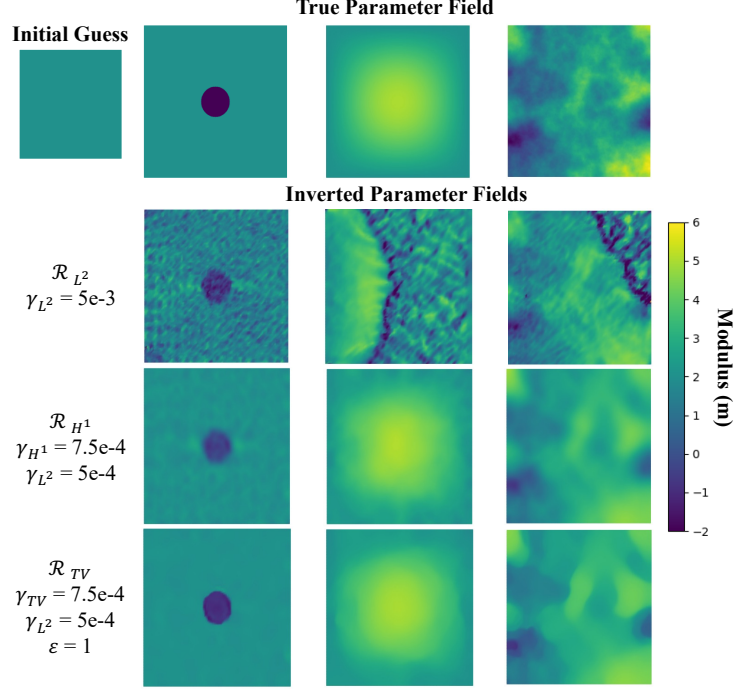


Fig. 5. Inverse results for Hyperelasticity for varying regularization. Each inversion is based on a simple tensile experiment with $\sim 5\%$ strain. 10% and 5% noise were applied to the image brightness and force measurements, respectively. The inverse problem is solved using a mesh size of 100×100 and a speckle correlation length of 0.01.

3.4. Mesh-independent Performance

The resolution of the mesh must be chosen to sufficiently be able to resolve the physics (e.g., feature size, such as for the thin, branched crack in Figure 4). We show inversions for 40,401, 160,801, and 361,201 parameter values (degrees of freedom) relating to 100×100 , 200×200 and 300×300 meshes, respectively. Previous IDIC and FEMU studies handled parameter fields in low dimensions such as two [6], three [7] and eight [64] scalar values. Increasing the number of degrees of freedom leads to smaller features being captured. Additionally, fine mesh resolution is important for representing the image speckle pattern. As we refine the mesh, the speckle pattern is better represented which leads to a better inversion as seen in Figure 6. The synthetic images were generated on a 500×500 mesh and then corrupted with noise; Since, the ∞ -IDIC was conducted on coarser meshes there is some information loss in the finite element representation of the images. Refining the mesh is necessary to minimize this effect, but it comes at additional computational costs.

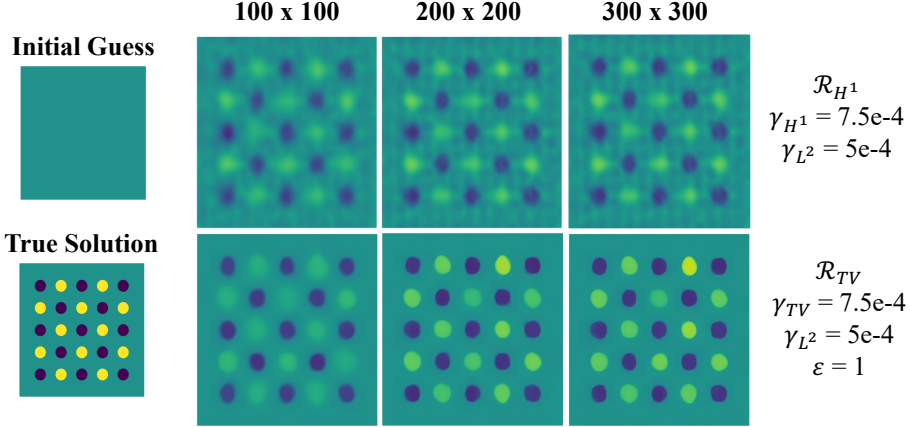


Fig. 6. Performance of ∞ -IDIC with H^1 and TV regularization for the sample problem with varying mesh resolution. Each inversion is based on a simple tensile experiment with $\sim 5\%$ strain. 10% and 5% noise were applied to the image brightness and force measurements, respectively. The speckle correlation length is 0.01.

A key benefit of formulating the problem in infinite dimensions is mesh-independent optimization behavior, leading to a scalable inversion algorithm. Refining the mesh only increases the cost per PDE solve and not the number of optimization iterations. Figure 7 demonstrates the mesh-independent performance of ∞ -IDIC for H^1 regularization on the parameter where the mesh increased from 100×100 , 200×200 to 300×300 . We show the relative cost (cost at each Newton iteration normalized by the initial cost) and the gradient norm. After 30 iterations, the relative cost plateaus independent of mesh size. The gradient norm exhibits similar convergence behavior, but the refined mesh leads to more jumps in the gradient values.

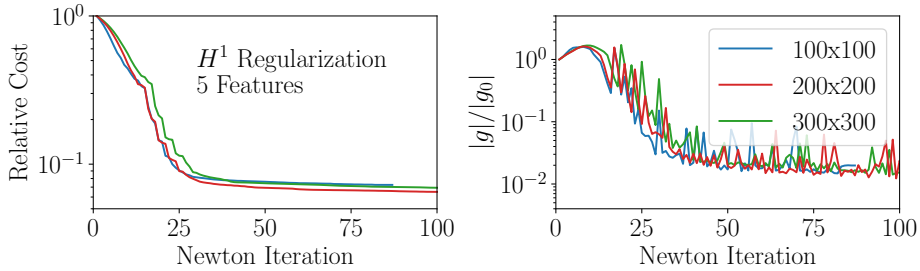


Fig. 7. Mesh-independent performance of ∞ -IDIC for H^1 regularization on the parameter where the mesh increased from 100×100 , 200×200 to 300×300 . This is a tensile experiment using a 5×5 features on hyperelastic medium with $\sim 5\%$ strain. 10% and 5% noise were applied to the image brightness and force measurements, respectively. The speckle correlation length is 0.01.

We demonstrate mesh-independent performance using the primal-dual TV formulation in Figure 8. It is clear from the results that ∞ -IDIC performs independently of the mesh size. The relative cost reduction and normalized gradient norm demonstrate similar convergence behavior among the varying mesh sizes. However, varying the number of features in m_{true} results in different convergence behavior, as the complexity of the problem increases. With more features, the problem requires more iterations to converge because each experiment starts from the same initial guess of $m(x) = 2$, and this guess is further from the true solution when more features are present. Additionally, the finer mesh better represents the image speckle, an effect that may be more pronounced for more complex problems with smaller features (e.g., 7×7). This effect may also be influenced by the regularization weighting. The tuning of the regularization parameter likely depends on the specific problem, and all examples in Figure 8 use the same regularization weighting. Ultimately, each experiment converges in a mesh-independent manner leading to a scalable algorithm that can be refined to represent the image speckle and resolve finer material features.

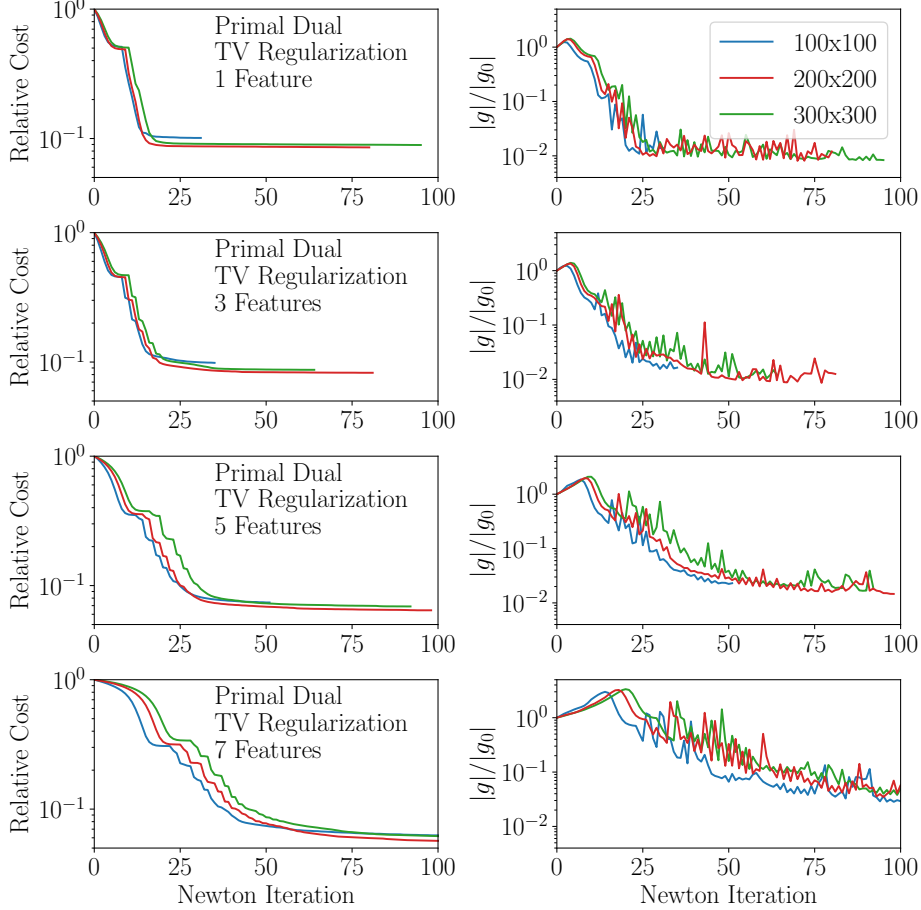


Fig. 8. Mesh-independent performance of ∞ -IDIC for Hyperelasticity when using the primal-dual total variation formulation ($\gamma_{L^2} = 5 \times 10^{-6}$, $\gamma_{TV} = 7.5 \times 10^{-4}$) where the mesh increased from 100×100 , 200×200 to 300×300 .

Mesh Size	Wall Clock Time (mins)
100×100	2.20
200×200	7.30
300×300	35.26

Table 2. Timing results for the three different mesh sizes. ∞ -IDIC was run on an M2 Macbook Pro with 16 Gb of memory in serial.

Increasing the mesh increases the computation time due to the PDE solves being more expensive. Table 2 shows the wall clock time for the three different mesh sizes. While the optimization behavior remains consistent, the computational cost increases significantly with the mesh size. The cost of the PDE solve and calculating $I(x + u(x))$ and $\nabla I(x + u(x))$ are the primary drivers of the computational cost. Notably, Pan et al. show that common DIC software (Vic-2D) took from 7 to 50 minutes to solve for the displacement field of a tensile coupon [65]. ∞ -IDIC is competitive with these times regardless of the solving a more complex inverse problem.

3.5. ∞ -IDIC Robustness

In this section, we investigate the robustness of ∞ -IDIC to noise in the images and error in the force measurement. Subsequently, we discuss how the speckle size impacts the resolution of inferred features. This leads us to show how the forcing condition also influences the inversion. To further investigate the nuances

of ∞ -IDIC, we introduce an expected information gain (EIG) heuristic that provides comparison between inversions. We use this to further investigate the influence of speckle size and magnitude of applied force.

3.5.1. Image Noise and Force Measurement Error

In real-world experiments observational noise and error are inevitable, so we investigate their influence on the inferred results. Appendix A explains synthetic data generation with noise corruption. We show in Figure 9 that with image brightness noise, ∞ -IDIC infers the correct solution as at 50% brightness noise and the inversion maintains 20 % accuracy. Previous studies show integrated DIC inverting for two parameters at 3% noise in the image brightness values begins to show significant error for single experiments, and requires multiple experiments to achieve the same accuracy [6]. On the other hand, for error in the force measurement, above 25% error the inversion fails to find the spatial distribution of the features. A notable advantage of ∞ -IDIC is its robustness to force measurement errors, which is particularly beneficial in contexts where measuring forces is challenging. For instance, when characterizing composite microstructure inside a Scanning Electron Microscope with IDIC, the force is applied outside the image field of view, necessitating an approximation of the local boundary conditions [10] such that being robust to force measurement is helpful. Regardless, we suspect that in most settings, the force measurement error is sufficiently less than 25%.

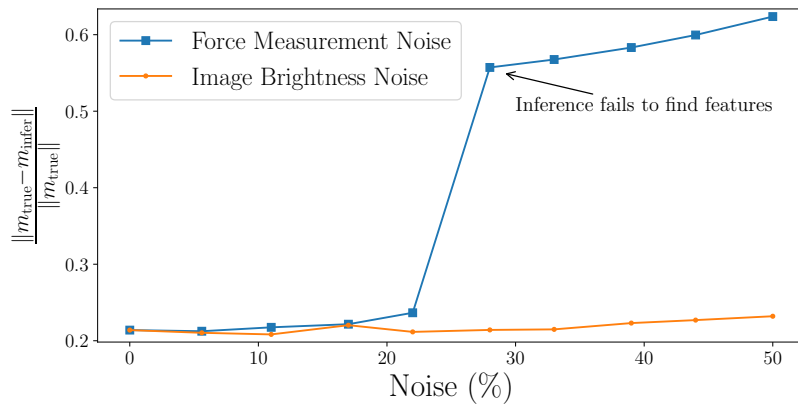


Fig. 9. Effect of noise on the accuracy of the inferred results for Hyperelasticity where m_{true} was a 5×5 case. Error in the force measurement eventually led to the inferred solutions not discovering the features, whereas image brightness noise had little effect.

3.5.2. Relationship Between Speckle Size and Resolution of Inferred Features

The speckle pattern is used during DIC to provide a contrast for comparison between the reference and deformed images. There are many methods to generate speckle patterns (e.g., spray paint, airbrush, stamping, etc.) and it is of interest to understand how the speckle pattern influences the inversion in high dimensions, particularly when the feature size is on similar length scales to the speckle size. As presented in the appendix, we generate the synthetic speckle by applying a threshold to a Gaussian random field with a tunable correlation length [63]. We therefore study the effect of speckle size by considering an inverse problem for a parameter field consisting of a single void in the domain, and compare the solutions of the inverse problems for various combinations of speckle and feature sizes.

Figure 10 shows the inverted parameter fields for the void, where the size of the void (feature) is reduced from $L/5$ to $L/50$ ($L = 1$ is the length of the domain) and the speckle correlations are reduced from 2.5×10^{-1} to 2.44×10^{-4} . Generally, we observe that smaller speckle sizes improve inversion results, allowing us to recover smaller features in the underlying parameter field. But, surprisingly, even when the speckle size is larger than the feature it is still discovered. We speculate that the edges of the speckle pattern (where ∇I_1 is largest) contain vital information for ∞ -IDIC, being the main contribution to the derivative of the misfit functional (see Equation (14)), such that large speckle sizes can be informative. Yet, we observe that when the feature size is reduced to $L/50$, larger speckle sizes fail to detect the void in the parameter field, while the smallest speckle size can still accurately recover the void.

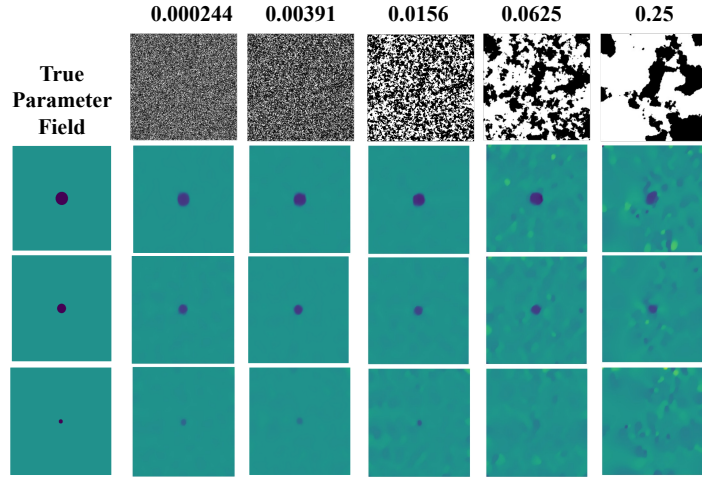
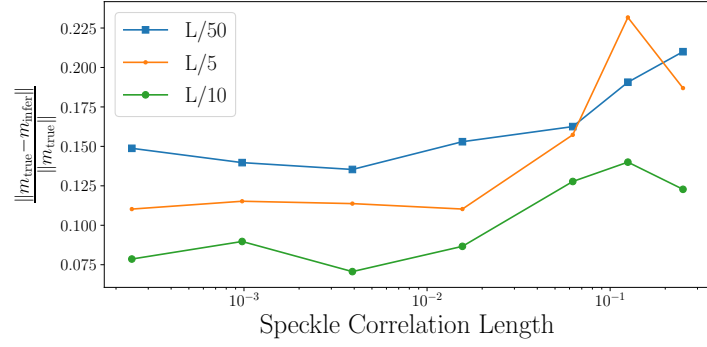


Fig. 10. The resolution of inverted features is influenced by the speckle size. ∞ -IDIC infers the correct solution even when the speckle size is sufficiently smaller than the feature size. 10% and 5% noise were applied to the image brightness and force measurements, respectively. The loading condition is a simple tensile experiment with $\sim 5\%$ strain.

3.5.3. Varying the Forcing Condition

In addition to the speckle pattern, the load configuration is a necessary component of the experiment. Figure 11 shows inversion solutions of the 3×3 grid modulus pattern using three different loading conditions; compression ($t_{\text{normal}} = -0.50$), tension ($t_{\text{normal}} = 0.50$), and bending ($t_{\text{shear}} = -0.25$). The errors compared to the true solution, $\|m_{\text{true}} - m_{\text{infer}}\| / \|m_{\text{true}}\|$, are 0.233, 0.219, and 0.218 for compression, tension and bending, respectively. We observe that the relative cost and gradient norm reduce more slowly for compression than the other loading conditions. Bending converges the quickest as shown in Figure 11 and appears to provide maximum information at the top of the domain where the sample is in tension. This aligns with the observation that tension is more informative than compression. Section 3.6.1 will consider coupling multiple experiments in a single inverse problem to further improve the inversion by providing more information.

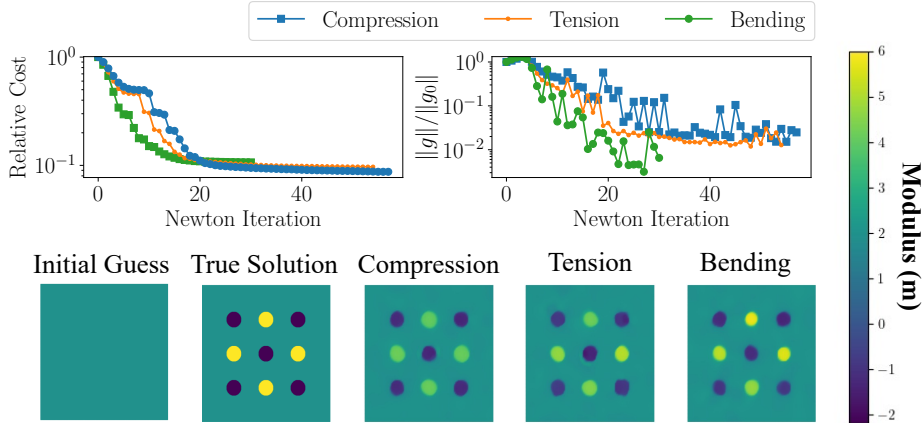


Fig. 11. ∞ -IDIC infers the correct solution for varying experiment boundary conditions. 10% and 5% noise were applied to the image brightness and force measurements, respectively. For each of the loading experiments, $\sim 5\%$ strain was applied. The inverse problem is solved using the primal-dual total variation formulation ($\gamma_{L^2} = 5 \times 10^{-6}$, $\gamma_{TV} = 7.5 \times 10^{-4}$) with a mesh size of 100×100 , and the speckle correlation length is 0.01.

3.5.4. An Information Gain Heuristic

We now introduce an expected information gain (EIG) metric from Bayesian optimal experimental design, as a means of evaluating how well-informed different experimental setups are relative to each other. EIG is a statistical quantity that measures how much information is gained about a random variable (in our case m) from the observation of another (possibly correlated) random variable (in our case the observation I_1). Given a means of computing EIG, we can make relative information comparisons by evaluating EIG for differing experimental setups (e.g., different traction conditions or different speckle patterns). While EIG formally requires a statistical inference framework for its derivation, in the case of linear inverse problems with Gaussian priors, it can be computed from a closed form expression [66]. The Gaussian prior plays a role analogous to the quadratic regularization (e.g., $H^1(\Omega)$ but not TV) in our presentation, so for this reason we restrict our attention to $H^1(\Omega)$ regularization in this section. In the context of linear Bayesian inverse problems, EIG is computed as follows. First, one computes the maximum a posteriori (MAP) point, which in this setting coincides with the solution of the deterministic inverse problem m^* . Second, one computes the generalized eigenpairs $\{(\lambda_i, \psi_i) | \lambda_i \geq \lambda_j \forall i < j\}_{i=1}^{\infty}$ at m^* , here written in variational form

$$\langle D_{mm} \Phi(u(m^*, t); I_0, I_1) \psi_i, \tilde{m} \rangle_{\mathcal{M}} = \lambda_i \langle \mathcal{C}^{-1} \psi_i, \tilde{m} \rangle_{\mathcal{M}} \quad \forall \tilde{m} \in \mathcal{M}, \quad (33)$$

where \mathcal{C}^{-1} is the self adjoint operator that induces the norm that is used in the regularization. For $H^1(\Omega)$ regularization this operator is the Laplace operator $\mathcal{C}^{-1} = -\gamma \Delta + \delta \mathcal{I}_{\mathcal{M}}$ with homogeneous Neumann boundary conditions, for appropriate choices of γ, δ . Then EIG, hereby denoted as Ψ is computed from the following formula

$$\Psi = \log \det \left(I_{\mathcal{M}} + \mathcal{C}^{1/2} D_{mm} \Phi \mathcal{C}^{1/2} \right) = \sum_{i=1}^{\infty} \log(1 + \lambda_i) \approx \sum_{i=1}^r \log(1 + \lambda_i), \quad (34)$$

where r is chosen such that $\lambda_j \approx 0$ for all $j > r$.

Remark 3.1. We note that the use of $H^1(\Omega)$ in the context of infinite-dimensional Bayesian inverse problems is abusive, as in this case \mathcal{C}^{-1} is not trace class for $\Omega \subset \mathbb{R}^d$ with $d > 1$; thus there is no meaningful limit of this problem. However, for a fixed discretization of the problem, this formulation leads to a meaningful approximation of EIG. For a more detailed discussion of EIG and the Bayesian interpretation of ∞ -IDIC see Appendix B.

In this interpretation of EIG, (34) suggests that the curvature of the Hessian located at the solution of the deterministic inverse problems indicates information gained in the solution of the optimization problem. Steeper basins around m^* indicate a well-informed solution of the inverse problem, while flatter basins indicate less informed solutions; this makes sense as in the latter case a perturbation of m^* would lead to a comparatively small change in Φ . Since we are equipped to compute all of these quantities via their

derivations in section 2.3, we propose using this linearized form of EIG as an information gain metric to compare trade offs in how the relative information gained between different choices of hyperparameter used in the inverse problem formulation. We note that this EIG metric can be efficiently approximated by the use of randomized generalized eigenvalue solvers [67, 68], since only the first r eigenpairs are needed as in (34). These eigenvalue solvers are implemented in `hippylib` [63].

3.5.5. Information Gain Increases with Refining Speckle

Now that we have defined a heuristic to compare between experiments, we run an experiment where the mesh, number of features, and regularization are held constant at 100×100 , 3×3 and $\gamma_{H^1} = 5 \times 10^{-4}$, while the speckle correlation length is refined. Figure 12 demonstrates that the EIG improves until plateauing around a speckle correlation length of 0.0156. This aligns with the observation in Figure 10 that the inversion improves with refining speckle size, until the speckle is sufficient for the feature size. Notably, the relative cost in Figure 12 also decays quicker with increasing speckle size suggesting that less Newton iterations are required. The relative cost, EIG, and inverted parameter fields are in agreement that refining the speckle size leads to improved inversions until a certain point where there is a diminishing return.

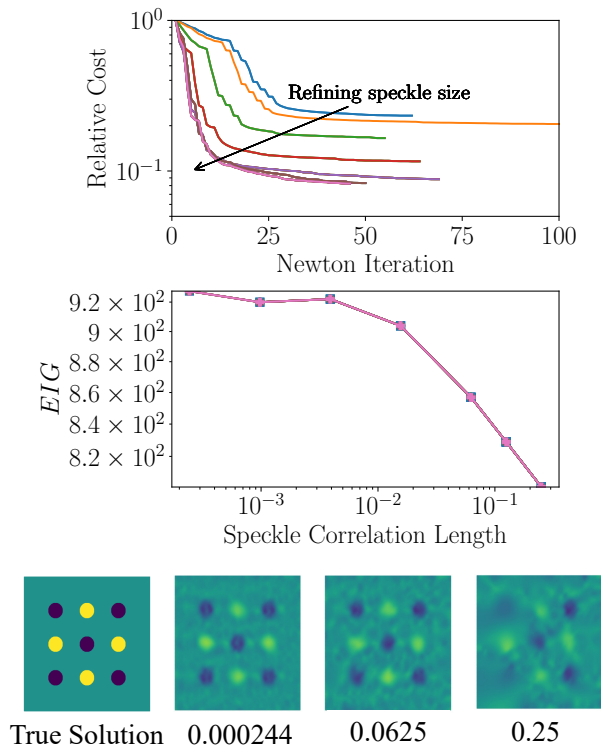


Fig. 12. The speckle size of the synthetic images is varied to demonstrate the relationship between the EIG and speckle size. The number of features is held constant at 3×3 while the speckle size is varied. Here, 0% noise is applied in order to not corrupt the eigenvalues. The loading condition is a simple tensile experiment with 5% strain. An H^1 regularization is used.

3.5.6. Information Gain Increases with Larger Forcing

Recall that an important aspect of DIC experiments to consider is the applied forces, as the loading condition determines the displacement observations. We again consider the 3×3 grid modulus pattern, but now we solve the inverse problems corresponding to various tensile loads from $t_{\text{normal}} = 0$ to $t_{\text{normal}} = 1$. As in the EIG analysis for speckle size, we use 0% noise for the synthetic images and adopt an H^1 regularization in the inversion. Notably, often in DIC, large deformations are difficult to handle due to ill-posedness and so a series of images to take smaller steps [69]. Yet, since IDIC benefits from incorporating the governing equations of the continuum mechanics we observe that a single pair of images can be used at large deformations ($\sim 10\%$) as seen in Figure 13. We plot the resulting EIG computed at the MAP points also in Figure 13. We observe that as the applied force increases, the EIG increases and ∞ -IDIC is able to resolve the modulus field

better. By increasing the deformation of the speckle pattern, ∞ -IDIC is provided with more informative data.

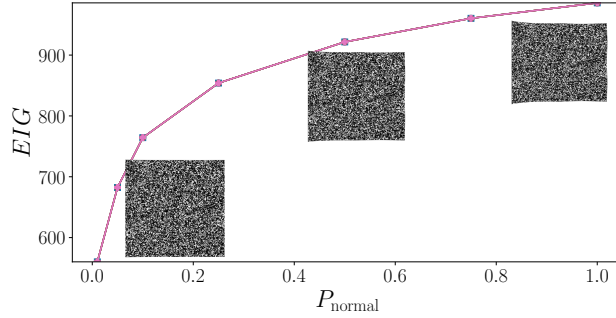


Fig. 13. The applied normal force, t_{normal} of the synthetic images is increased to demonstrate that additional deformation improves the inversion. The number of features is held constant at 3×3 while the speckle size is varied. Here, 0% noise is applied in order to not corrupt the eigenvalues. The loading condition is a simple tensile experiment with 5% strain. An H^1 regularization is used.

3.6. Useful Extensions of ∞ -IDIC

We will present two forward-looking aspects of ∞ -IDIC. First, we consider incorporating multiple experiments into a single inverse problem. Then, we show that we can discover stress concentrations by capturing the effects of the spatially-varying material properties.

3.6.1. Learning More Through Multiple Experiments

We showed that increasing the magnitude of the force improved the EIG (see section 3.5.6) and the inversion solution changed varying the loading configuration (see section 3.5.3). This section demonstrates the marriage of multiple experiments into a single inverse problem; It has been shown before that using multiple experiments improved the inversion of five isotropic elastic parameters in an IDIC formulation [7]. To formulate ∞ -IDIC with multiple experiments, we update the cost functional for ∞ -IDIC to be a sum over the experiments,

$$\hat{\mathcal{J}}(m) := \left(\frac{1}{N_{\text{exp}}} \sum_{i=1}^{N_{\text{exp}}} \Phi(u(m, t_i); I_{0,i}, I_{1,i}) \right) + \mathcal{R}(m), \quad (35)$$

where N_{exp} is the number of experiments, $(t_i)_{i=1}^{N_{\text{exp}}}$ are the different loading conditions, and $(I_{0,i})_{i=1}^{N_{\text{exp}}}$ and $(I_{1,i})_{i=1}^{N_{\text{exp}}}$ are the undeformed and deformed images corresponding to each experiment. Similarly, the forward, adjoint, and gradient equations become a sum over the experiments and subsequently the incremental forward, incremental adjoint, and incremental gradient as well. Including the $1/N_{\text{exp}}$ ensures that the regularization weighting is consistent with a single experiment.

We first show results for a single tension ($t_{\text{normal}} = 0.50$) experiment to contrast handling four experiments simultaneously: tension ($t_{\text{normal}} = 0.50$), compression ($t_{\text{normal}} = -0.50$), downwards bending ($t_{\text{shear}} = -0.10$), and upwards bending ($t_{\text{shear}} = 0.10$). The 3×3 feature example noticeably better recovers the true solution's modulus values with the additional experiments. Small features that were not recovered with a single experiment are improved such as the ears of the Texas Longhorn. Lastly, spurious features are removed upon with the addition of more experiments. This is especially the case with the Voronoi tessellations. We find that varying the loading conditions improves the inversion for ∞ -IDIC.

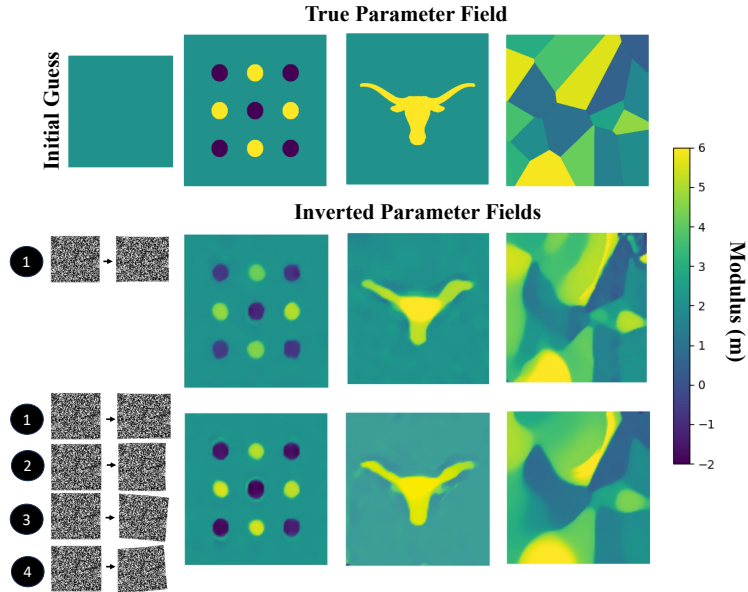


Fig. 14. ∞ -IDIC infers the correct solution for multiple experiments. 10% and 5% noise were applied to the image brightness and force measurements, respectively. For each of the loading experiments, $\sim 5\%$ strain was applied, albeit t_{normal} and t_{normal} varied. The inverse problem is solved using the primal-dual total variation formulation ($\gamma_{L^2} = 5e-6$, $\gamma_{TV} = 7.5e-4$) with a mesh size of 100×100 , and the speckle correlation length is 0.01.

3.6.2. Post-processing for Modulus-Induced Stress Concentrations

In DIC experiments, it is customary to post-process the displacement field to a strain field using the Green–Lagrange strain relationship,

$$\epsilon = \frac{1}{2} (\nabla u + \nabla u^T). \quad (36)$$

For instance, Ncorr, an open-source DIC software, provides this functionality [70]. The strains are visualized as spatially varying scalar fields by selecting one direction in the tensor, such as ϵ_{xx} which aligns with the directional of the normal applied load. We show results of ϵ_{xx} using linear elasticity shown in Figure 15. The strain fields provide insight into loading behavior. Yet, engineering systems are often designed to stress criteria. In the setting of heterogeneous materials, the stress field is *different* from the strain field. We use a common stress criterion for linear elasticity, the von Mises stress which is used to predict yielding in materials under uniaxial tension. The von Mises stress is a scalar field that represents the stress state of the material and is convenient for visualization. For linear elasticity, a von Mises stress field, σ_{vm} , can be calculated as

$$\sigma_{vm} = \sqrt{\frac{3}{2} s : s} \quad (37a)$$

$$s = \sigma - \frac{1}{3} \text{tr}(\sigma) I. \quad (37b)$$

We show post-processing results for linear elasticity only, but similar techniques could be done for the hyperelasticity case. Figure 15 shows the true and inverted strain and stress fields for the same features as in Figure 3. The first example involves a 3×3 grid of alternating voids stiff regions where the inferred strain and stress fields capture the general trends successfully. The strain field shows limited deformation in stiff interfaces and large deformation in the void regions while stress concentrations are observed around voids, indicating potential high-risk regions for failure. In the Voronoi tessellation, the strain field shows maximum strain in grains with lower modulus and minimal strain in high modulus regions. The inferred fields properly identify stress concentrates at grain interfaces, aligning with the true solution. Lastly, with the isotropic Gaussian random field there are stress concentrations properly identified around high modulus

regions. ∞ -IDIC captures general trends in strain and stress fields even for a single experiment at low strains ($\sim 0.2\%$), leading to potential applications in non-destructive failure predictions.

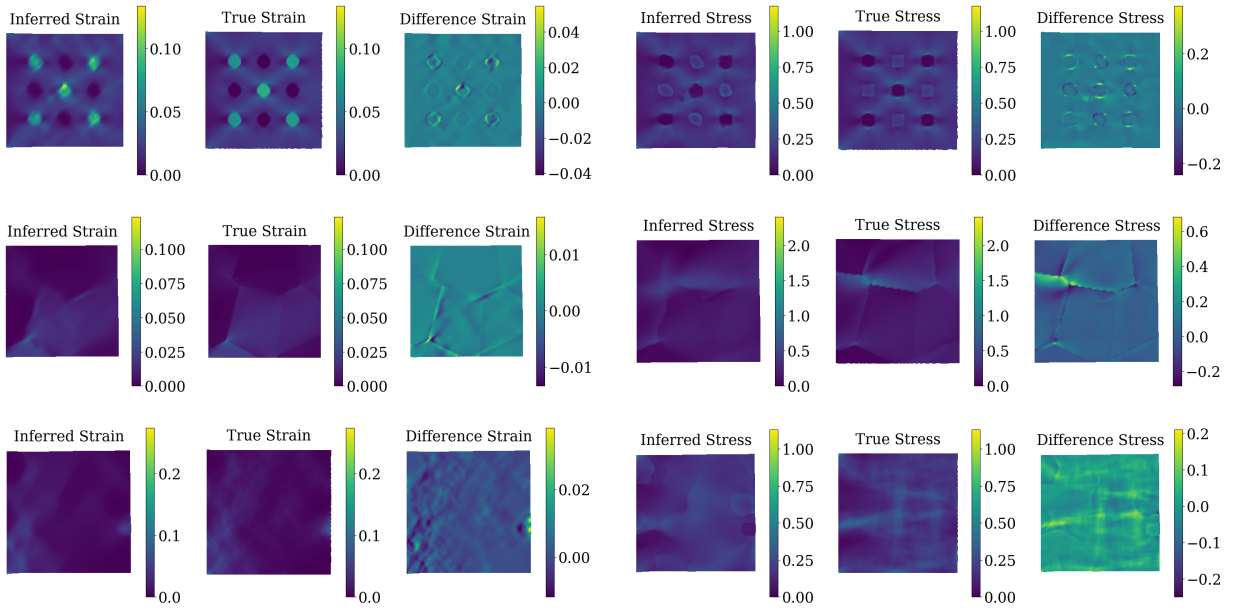


Fig. 15. Strain and stress fields for the true and inferred solutions for linear elasticity. The first row shows the strain fields for the 3×3 grid of alternating voids and stiff regions. The second row shows the strain fields for the Voronoi tessellation. The third row shows the strain fields for the isotropic Gaussian random field.

4. Conclusion

In summary, we present ∞ -IDIC, a scalable and efficient framework for the inference of spatially-varying heterogeneous materials that parametrize continuum mechanical models. The goal of ∞ -IDIC is the characterization of damage, defects, anomalies and inclusions in material specimen through nondestructive evaluation methods, without any prior knowledge of their morphology or strength. These effects cannot be directly observed, but must be reconstructed indirectly through inverse problems. Since heterogeneous material fields are formally infinite dimensional, we require (i) rich information in the form of observational data to faithfully reconstruct the material properties, and (ii) an infinite-dimensional formulation of the inverse problem. To address the first concern we build on IDIC, which like traditional DIC utilizes images of material deformation as observational data, but unlike DIC monolithically couples the image registration to the governing equations of equilibrium to enforce physical consistency. To address the second concern we formulate a general formulation in a function space setting to allow for the inference of arbitrary material parameter fields. The main challenges of this approach are due to the mathematical and computational challenges posed by the infinite-dimensional formulation, namely ill-posedness and high dimensionality; addressing these challenges is a focus of this work. First, to tackle the inherent ill-posedness we consider various regularization schemes, namely $H^1(\Omega)$ and total variation for the inference of smooth and sharp features, respectively. To address the computational costs associated with the discretized problem, we utilize an cost efficient and dimension-independent inexact-Newton CG framework for the solution of the regularized inverse problem.

In the numerical results, we demonstrate that ∞ -IDIC can recover spatially varying material properties with remarkable accuracy and sharpness using examples of linear elasticity and hyperelasticity. We demonstrate the ability to invert for 40,401, 160,801, and 361,201 parameter values while maintaining mesh-independent behavior. We conduct numerical experiments to investigate how robust ∞ -IDIC is to noise in the image and force data, variation in speckle sizes, and varying force conditions. We find that refining the speckle size and increasing the force improves inversions. This leads us to show that coupling multiple experiments in a single ∞ -IDIC inverse problem further improves inversions. Lastly, we demonstrate that

along with the inversion of the spatially varying features, we can invert for stress concentrations arising from heterogeneity which may potentially be useful for non-destructive evaluation.

We recognize that there is a need to characterize complex, heterogeneous fields such as welds, composites, tissues, etc. We believe that ∞ -IDIC can be a powerful tool for such challenges due to its ability to invert for sharp features and its mesh-independent behavior. We hope that this enables researchers to better characterize complex materials and ultimately build more accurate models. The ∞ -IDIC framework can be extended to other settings such as fluid mechanics where an image registration problem can be coupled with a PDE model (e.g., using particle image velocimetry observational data). In this work we have focused on 2D PDE problems with simple geometries for the sake of a methodological exposition. Extending ∞ -IDIC to complex geometries, three-dimensional and time-dependent problems which would be a significant step forward for material characterization, and constitutes future work.

5. Acknowledgments

The authors are grateful to Umberto Villa for insightful discussions, and help with software questions related to `hIPPYlib` [63]. The authors are also grateful for Graham Pash for insightful discussions.

References

- [1] F. Pierron, *Material Testing 2.0: A brief review*, 2023.
- [2] F. Pierron, M. Grédiac, *Towards Material Testing 2.0. A review of test design for identification of constitutive parameters from full-field measurements*, *Strain* 57 (2021).
- [3] J. Blaber, B. Adair, A. Antoniou, *Ncorr: Open-Source 2D Digital Image Correlation Matlab Software*, *Experimental Mechanics* 55 (2015) 1105–1122.
- [4] G. Besnard, F. Hild, S. Roux, "Finite-element" displacement fields analysis from digital images: Application to Portevin-Le Châtelier bands, *Experimental Mechanics* 46 (2006) 789–803.
- [5] V. Kosin, A. Fau, C. Jailin, F. Hild, T. Wick, *Parameter identification of a phase-field fracture model using integrated digital image correlation*, *Computer Methods in Applied Mechanics and Engineering* 420 (2024).
- [6] F. Mathieu, H. Leclerc, F. Hild, S. Roux, *Estimation of Elastoplastic Parameters via Weighted FEMU and Integrated-DIC*, *Experimental Mechanics* 55 (2015) 105–119.
- [7] J. Neggers, F. Mathieu, F. Hild, S. Roux, *Simultaneous full-field multi-experiment identification*, *Mechanics of Materials* 133 (2019) 71–84.
- [8] Y. Li, J. Zhao, J. Zhou, Y. Yang, X. Huang, Z. Liu, *Local-Micro-Zone-Wise Time-Resolved Integrated Digital Image Correlation for Evaluating the Mechanical Properties of Welding Joints*, *Experimental Techniques* (2022).
- [9] S. B. Lindström, H. Wemming, Z. Kapidžić, M. S. Loukil, M. Segersäll, *Integrated digital image correlation for mechanical characterization of carbon fiber-reinforced polymer plates*, *Composite Structures* 305 (2023).
- [10] O. Rokoš, R. H. Peerlings, J. P. Hoefnagels, M. G. Geers, *Integrated digital image correlation for micro-mechanical parameter identification in multiscale experiments*, *International Journal of Solids and Structures* 267 (2023).
- [11] K. T. Kavanaght, W. Clough, *Finite Element Applications In The Characterization Of Elastic Solids*, Technical Report, 1971.
- [12] Y. Sun, J. H. L. Pang, C. Khuen Wong, F. Su, *Finite element formulation for a digital image correlation method*, Technical Report, 2005.
- [13] P. Lenny, A. S. Caro-Bretelle, E. Pagnacco, *Identification from measurements of mechanical fields by finite element model updating strategies: A review*, *European Journal of Computational Mechanics* 18 (2009) 353–376.
- [14] H. Waisman, E. Chatzi, A. W. Smyth, *Detection and quantification of flaws in structures by the extended finite element method and genetic algorithms*, *International Journal for Numerical Methods in Engineering* 82 (2010) 303–328.
- [15] A. Elouneq, D. Sutula, J. Chambert, A. Lejeune, S. P. Bordas, E. Jacquet, *An open-source fenics-based framework for hyperelastic parameter estimation from noisy full-field data: Application to heterogeneous soft tissues*, *Computers and Structures* 255 (2021).
- [16] S. Ereiz, I. Duvnjak, J. Fernando Jiménez-Alonso, *Review of finite element model updating methods for structural applications*, 2022.
- [17] S. S. Fayad, E. M. Jones, D. T. Seidl, P. L. Reu, J. Lambros, *On the Importance of Direct-Levelling for Constitutive Material Model Calibration using Digital Image Correlation and Finite Element Model Updating*, *Experimental Mechanics* 63 (2023) 467–484.
- [18] A. P. Ruybalid, J. P. Hoefnagels, O. van der Sluis, M. G. Geers, *Comparison of the identification performance of conventional FEM updating and integrated DIC*, *International Journal for Numerical Methods in Engineering* 106 (2016) 298–320.
- [19] R. B. Lehoucq, P. L. Reu, D. Z. Turner, *The Effect of the Ill-posed Problem on Quantitative Error Assessment in Digital Image Correlation*, *Experimental Mechanics* 61 (2021) 609–621.
- [20] A. R. Skovoroda, L. Lubinski, S. Y. Emelianov, M. O'Donnell, *Reconstructive elasticity imaging for large deformations*, *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* 46 (1999) 523–535.
- [21] Y. Zhu, T. J. Hall, J. Jiang, *A finite-element approach for young's modulus reconstruction*, *IEEE transactions on medical imaging* 22 (2003) 890–901.
- [22] Y. Liu, L. Sun, G. Wang, *Tomography-based 3-d anisotropic elastography using boundary measurements*, *IEEE Transactions on Medical Imaging* 24 (2005) 1323–1333.
- [23] S. Goenezen, P. Barbone, A. A. Oberai, *Solution of the nonlinear elasticity imaging inverse problem: The incompressible case*, *Computer methods in applied mechanics and engineering* 200 (2011) 1406–1420.
- [24] Y. Mei, R. Fulmer, V. Raja, S. Wang, S. Goenezen, *Estimating the non-homogeneous elastic modulus distribution from surface deformations*, *International Journal of Solids and Structures* 83 (2016) 73–80.
- [25] Y. Mei, B. Stover, N. A. Kazerooni, A. Srinivasa, M. Hajhashemkhani, M. Hematiyan, S. Goenezen, *A comparative study of two constitutive models within an inverse approach to determine the spatial stiffness distribution in soft materials*, *International journal of mechanical sciences* 140 (2018) 446–454.
- [26] H. W. Engl, M. Hanke, A. Neubauer, *Regularization of inverse problems*, volume 375, Springer Science & Business Media, 1996.
- [27] C. R. Vogel, *Computational methods for inverse problems*, SIAM, 2002.
- [28] A. Tarantola, *Inverse problem theory and methods for model parameter estimation*, Society for Industrial and Applied Mathematics, 2005.
- [29] A. M. Stuart, *Inverse problems: A bayesian perspective*, 2010.
- [30] O. Ghattas, K. Willcox, *Learning physics-based models from data: perspectives from inverse problems and model reduction*, 2021.
- [31] H. G. Maas, A. Gruen, D. Papantoniou, *Particle tracking velocimetry in three-dimensional flows Part I. Photogrammetric determination of particle coordinates*, Technical Report, 1993.
- [32] N. A. Malik, T. Dracos, D. A. Papantoniou, *Particle tracking velocimetry in three-dimensional flows Part II: Particle tracking*, Technical Report, 1993.
- [33] M. R. Abdulwahab¹, Y. H. Ali, F. J. Habeeb, A. A. Borhana, A. M. Abdelrhman, S. M. A. Al-Obaidi, *A Review in Particle*

- Image Velocimetry Techniques (Developments and Applications), *Journal of Advanced Research in Fluid Mechanics and Thermal Sciences Journal homepage* 65 (2020) 213–229.
- [34] W. Brevis, Y. Niño, G. H. Jirka, Integrating cross-correlation and relaxation algorithms for particle tracking velocimetry, *Experiments in Fluids* 50 (2011) 135–147.
- [35] Y. G. Guezennec, R. S. Brodkey, N. Trigui, J. C. Kent, Algorithms for fully automated three-dimensional particle tracking velocimetry, Technical Report, 1994.
- [36] A. Schröder, D. Schanz, Annual Review of Fluid Mechanics 3D Lagrangian Particle Tracking in Fluid Mechanics, *Annu. Rev. Fluid Mech.* 2023 55 (2022) 511–540.
- [37] A. Ahmadzadegan, H. Mitra, P. P. Vlachos, A. M. Ardekani, Particle Image micro-Rheology (PIR) using displacement probability density function, *Journal of Rheology* 67 (2023) 823.
- [38] M. V. Wyawahare, P. M. Patil, H. K. Abhyankar, Image Registration Techniques: An overview, Technical Report 3, 2009.
- [39] D. L. G Hill, P. G. Batchelor, M. Holden, D. J. Hawkes, Physics in Medicine & Biology Medical image registration Medical image registration, Technical Report, 2001.
- [40] F. P. Oliveira, J. M. R. Tavares, Medical image registration: A review, *Computer Methods in Biomechanics and Biomedical Engineering* 17 (2014) 73–93.
- [41] Y. Fu, Y. Lei, T. Wang, W. J. Curran, T. Liu, X. Yang, Deep learning in medical image registration: A review, 2020.
- [42] A. Sotiras, C. Davatzikos, N. Paragios, Deformable medical image registration: A survey, *IEEE Transactions on Medical Imaging* 32 (2013) 1153–1190.
- [43] B. Rigaud, A. Simon, J. Castelli, C. Lafond, O. Acosta, P. Haigron, G. Cazoulat, R. de Crevoisier, Deformable image registration for radiation therapy: principle, methods, applications and evaluation, 2019.
- [44] R. Sandhu, S. Dambreville, A. Tannenbaum, Point set registration via particle filtering and stochastic dynamics, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32 (2010) 1459–1473.
- [45] J. Ma, J. Zhao, A. L. Yuille, Non-rigid point set registration by preserving global and local structures, *IEEE Transactions on Image Processing* 25 (2016) 53–64.
- [46] X. Yuan, A. Maharjan, Non-rigid point set registration: recent trends and challenges, *Artificial Intelligence Review* 56 (2023) 4859–4891.
- [47] B. Maiseli, Y. Gu, H. Gao, Recent developments and trends in point set registration methods, 2017.
- [48] Z. Min, A. Zhang, Z. Zhang, J. Wang, S. Song, H. Ren, M. Q. Meng, 3-D Rigid Point Set Registration for Computer-Assisted Orthopedic Surgery (CAOS): A Review From the Algorithmic Perspective, *IEEE Transactions on Medical Robotics and Bionics* 5 (2023) 156–169.
- [49] T. Liu, A. Merat, M. H. Makhmalbaf, C. Fajardo, P. Merati, Comparison between optical flow and cross-correlation methods for extraction of velocity fields from particle images, *Experiments in Fluids* 56 (2015).
- [50] B. Wang, Z. Cai, L. Shen, T. Liu, An analysis of physics-based optical flow, *Journal of Computational and Applied Mathematics* 276 (2015) 62–80.
- [51] S. T. H. Shah, X. Xuezi, Traditional and modern strategies for optical flow: an investigation, 2021.
- [52] R. Becker, B. Vexler, Optimal control of the convection-diffusion equation using stabilized finite element methods, *Numerische Mathematik* 106 (2007) 349–367.
- [53] M. Braack, Optimal control in fluid mechanics by finite elements with symmetric stabilization, *SIAM Journal on Control and Optimization* 48 (2009) 672–687.
- [54] M. Heinkenschloss, Mesh independence for nonlinear least squares problems with norm constraints*, 1993.
- [55] V. Akçelik, G. Biros, O. Ghattas, J. Hill, D. Keyes, B. van Bloemen Waanders, Parallel processing for scientific computing, chapter parallel algorithms for pde-constrained optimization, 2006.
- [56] A. Borzi, V. Schulz, Computational optimization of systems governed by partial differential equations, SIAM, 2011.
- [57] S. C. Eisenstat, H. F. Walker, Globally convergent inexact newton methods, *SIAM Journal on Optimization* 4 (1994) 393–422.
- [58] T. F. Chan, G. H. Golub, S. J. C. Sci, A nonlinear primal-dual method for total variation-based image restoration*, *SIAM J. SCI.COMPUT* 20 (1999) 1964–1977.
- [59] T. F. Chan, P. Mulet, On the convergence of the lagged diffusivity fixed point method in total variation image restoration, *SIAM Journal on Numerical Analysis* 36 (1999) 354–367.
- [60] D. Arnold, R. Falk, R. Winther, Finite element exterior calculus: from hodge theory to numerical stability, *Bulletin of the American mathematical society* 47 (2010) 281–354.
- [61] M. Herrmann, R. Herzog, S. Schmidt, J. Vidal-Núñez, G. Wachsmuth, Discrete total variation with finite elements and applications to imaging, *Journal of Mathematical Imaging and Vision* 61 (2019) 411–431.
- [62] M. S. Alnaes, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, G. N. Wells, The fenics project version 1.5, 2015.
- [63] U. Villa, N. Petra, O. Ghattas, Hippylib: An extensible software framework for large-scale inverse problems governed by pdes: Part i: Deterministic inversion and linearized bayesian inference, *ACM Transactions on Mathematical Software* 47 (2021).
- [64] R. Gras, H. Leclerc, F. Hild, S. Roux, J. Schneider, Identification of a set of macroscopic elastic parameters in a 3D woven composite: Uncertainty analysis and regularization, *International Journal of Solids and Structures* 55 (2015) 2–16.
- [65] B. Pan, K. Li, A fast digital image correlation method for deformation measurement, *Optics and Lasers in Engineering* 49 (2011) 841–847.
- [66] A. Alexanderian, P. J. Gloor, O. Ghattas, On Bayesian A-and D-optimal experimental designs in infinite dimensions, *Bayesian Analysis* 11 (2016) 671–695.
- [67] N. Halko, P.-G. Martinsson, J. A. Tropp, Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions, *SIAM review* 53 (2011) 217–288.
- [68] P.-G. Martinsson, J. A. Tropp, Randomized numerical linear algebra: Foundations and algorithms, *Acta Numerica* 29 (2020) 403–572.

- [69] B. Grant, H. Stone, P. Withers, M. Preuss, High-temperature strain field measurement using digital image correlation, *The Journal of Strain Analysis for Engineering Design* 44 (2009) 263–271.
- [70] J. Blaber, . B. Adair, . A. Antoniou, Ncorr: Open-Source 2D Digital Image Correlation Matlab Software (2014).
- [71] Z. Yao, Z. Hu, J. Li, A tv-gaussian prior for infinite-dimensional bayesian inverse problems and its numerical implementations, *Inverse Problems* 32 (2016) 075006.
- [72] D. Lv, Q. Zhou, J. K. Choi, J. Li, X. Zhang, Nonlocal tv-gaussian prior for bayesian inverse problems with applications to limited ct reconstruction., *Inverse Problems & Imaging* 14 (2020).
- [73] F. Lindgren, H. Rue, J. Lindström, An explicit link between gaussian fields and gaussian markov random fields: the stochastic partial differential equation approach, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73 (2011) 423–498.
- [74] U. Villa, T. O’Leary-Roseberry, A note on the relationship between PDE-based precision operators and Matérn covariances (2024).

Appendix A. Details on Numerical Implementation

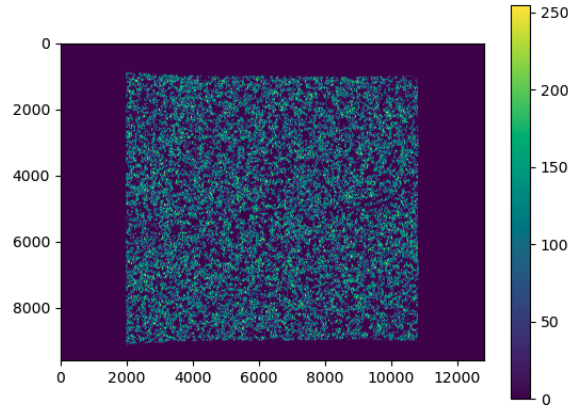


Fig. A.16. Image brightness noise at 50% added to the synthetically generated images. The noise is added to the images to simulate the effect of varying lighting conditions or general imaging errors in the experiment.

This section describes the process of generating synthetic data for the inverse problem. The assumption of a prior agreement between our understanding of the physics and the observed data leads us to commit what is known as an "inverse crime." Each numerical result is derived using the same PDE that is subsequently utilized to inform the inverse problem. Synthetic data is generated by solving the PDEs for the forward problem and saving speckled images before and after deformation. The speckled images are then used as the observational data in the subsequent inverse problem. In an actual experiment, the images would be captured using a camera, or similar sensor, while the speckled specimen is loaded. The following steps detail how we synthetically form speckled image data:

1. **Mesh construction:** Construct a mesh of size 500×500 where 500 is the number of elements in the mesh. The mesh is used to approximate Ω_0 via linear Lagrangian triangular elements to form finite element spaces $\mathcal{M}^h \in \mathcal{M}$ and $\mathcal{V}^h \in \mathcal{V}$. The parameter and state space are 1,002,001 and 502,002 dimensions, respectively.
2. **True parameter generation:** Generate a true field, $m_{\text{true}}(x)$, that represents the true material model. This field is used to generate the Young’s modulus, $E(x) = e^{m_{\text{true}}(x)}$, and the Poisson’s ratio $\nu(x) = 0.35$. The true field is generated through a handful of strategies geared towards representing complex material fields.
3. **Boundary conditions definition:** Define the traction condition, t , on the right boundary, Γ_R , by determining the forces, P , and the Dirichlet boundary condition, $u = 0$, on the left boundary, Γ_L . The Neumann boundary condition is defined as $t = 0$ on the top and bottom boundaries, Γ_T and Γ_B . Figure 2 illustrates the boundary conditions.
4. **Forward PDE solve:** Solve the forward PDE for the true material model to obtain the true displacement field, $u_{\text{true}}(x)$, and the true stress field, $\sigma_{\text{true}}(x)$.

5. **Image generation:** Generate a speckled image existing in the discretized parameter space, \mathcal{M} by taking a sample, i , from a bilaplacian prior. The correlation length of the bilaplacian prior, ℓ , regulates the respective correlation length of the speckles. As the speckles in DIC are black (0) or white (255), thresholding is applied. This is done by projecting i to a function space, \mathcal{M}^h , and transforming for the image, $I(x)$, using the following equation:

$$I(x) = 255 - 255 \tanh(100i(x) + 1). \quad (\text{A.1})$$

6. **Image mapping in deformed states & Brightness Error:** Map the simulated speckle to both the reference and deformed states, thereby generating the reference image (I_0) and the deformed image (I_1). A modified version of the FEniCS plotting tool operating in displacement mode is used. Instead of representing displacement, the color map represents the grayscale speckle values. The displaced image is corrupted with white noise, i.e., $N \sim \mathcal{N}(0, \sigma^2)$, where sigma is the noise level multiplied by the maximum pixel value to simulate image brightness error. Notably, in practice, the material domain, Ω_0 , is smaller than the image domain, Ω_I . To simulate this, the images are saved as .png files with additional white space around the material domain.
7. **Force measurement error simulation:** To mimic errors in force measurement data, corrupt the traction condition. For instance, if the true force value is $t_{\text{normal}} = 10$, the corrupted value is $t_{\text{normal}} = 9.5$ given a 5% error in a tensile measurement. The corrupted value is saved to a .txt file for subsequent use in the inverse problem.

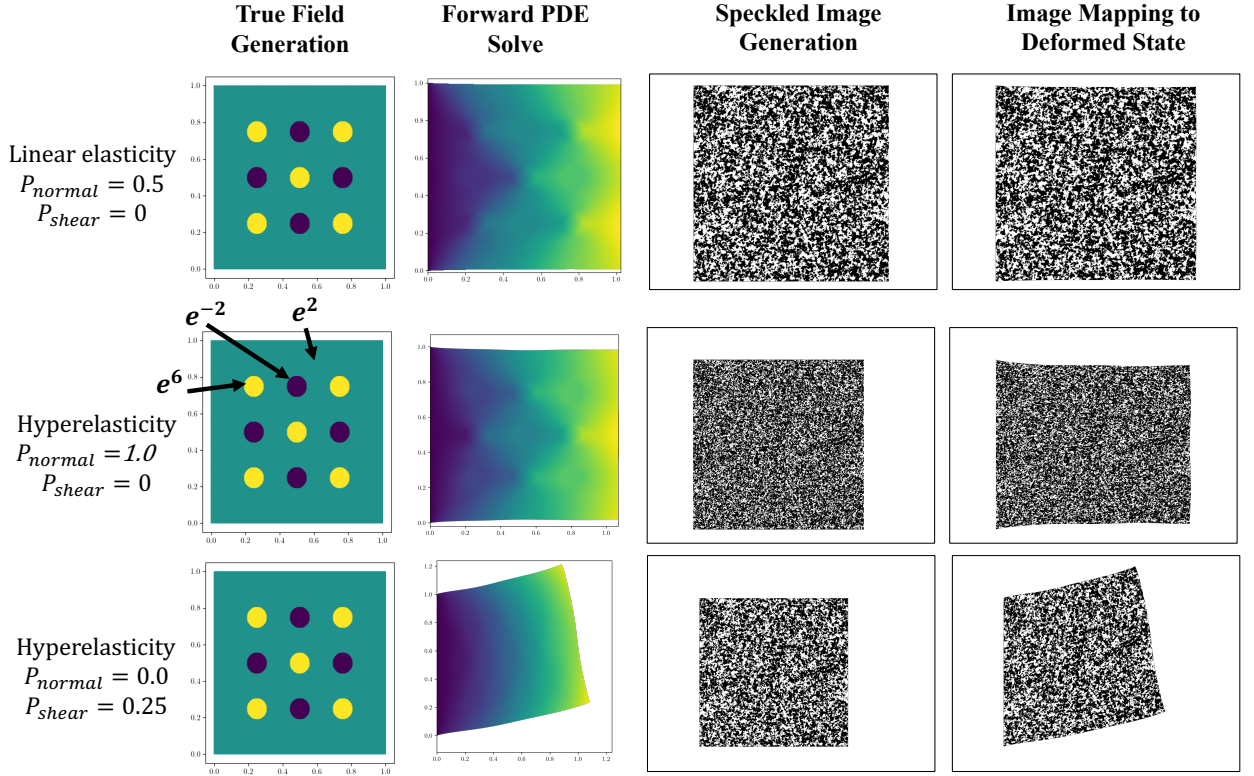


Fig. A.17. Schematic of the synthetic data generation process.

Appendix B. Bayesian ∞ -IDIC and expected information gain

We here present a Bayesian extension of ∞ -IDIC, in order to give sufficient background for the expected information gain (EIG) numerical results presented in 3.5.4. First we assume \mathcal{M} to be a separable Hilbert space, so basis representations are countable. We denote by $(\mathcal{M}, \mathcal{B}(\mathcal{M}))$, the measurable space with $\mathcal{B}(\mathcal{M})$

being the Borel σ -algebra generated by open sets in \mathcal{M} . Next we let $\mathcal{C} : \mathcal{M} \rightarrow \mathcal{M}$ be a trace class, self adjoint linear operator on \mathcal{M} , such that

$$\mathrm{tr}(\mathcal{C}) = \sum_{i=1}^{\infty} \langle \psi_i, \mathcal{C} \psi_i \rangle_{\mathcal{M}} < \infty, \quad \mathcal{C} = \mathcal{C}^*, \quad (\text{B.1})$$

where $\{\psi_i\}_{i=1}^{\infty}$ is any orthonormal basis for \mathcal{M} . The operator \mathcal{C} can induce an inner product and thereby a norm. Regularization schemes that are quadratics of the norm $\|\cdot\|_{\mathcal{C}^{-1}}^2$ induce Gaussian measures in infinite dimensional settings. The associated regularization

$$\mathcal{R}_{\mathcal{C}^{-1}}(m) = \|m - m_0\|_{\mathcal{C}^{-1}}^2 = \|\mathcal{C}^{-1/2}(m - \bar{m})\|_{L^2(\Omega)}^2, \quad (\text{B.2})$$

play the same role as a prior distribution, $\mu_{\text{prior}} = \mathcal{N}(m_0, \mathcal{C})$, in the Bayesian setting. Here $\mathcal{C}^{-1/2}$ denotes the symmetric square root of the operator \mathcal{C}^{-1} .

What is left to work towards Bayes is a likelihood measure; towards this end, we assume an additive Gaussian noise model for the observed images, e.g., where noise $\xi \sim \mathcal{N}(0, \mathcal{C}_{\text{noise}})$, where $\mathcal{C}_{\text{noise}} : \mathcal{I} \rightarrow \mathcal{I}$ is a self adjoint trace class linear operator. In this case our image misfit takes the following form

$$\Phi(u(m); I_0, I_1) = \left\| \mathcal{C}_{\text{noise}}^{-\frac{1}{2}} (I_1(x + u(x)) - I_0(x)) \right\|_{L^2(\Omega)}^2 \quad (\text{B.3})$$

and leads to the following likelihood

$$\pi_{\text{like}}(I_1|m) \propto \exp(-\Phi(u(m); I_0, I_1)). \quad (\text{B.4})$$

Bayes Theorem provides a rigorous framework for statistically inferring m as a *posterior distribution*, $\mu_{\text{post}}^{I_1} = \mu_{\text{post}}(m|I_1)$ given the observational data I_1 through the lens of the PDE model:

$$\frac{d\mu_{\text{post}}^{I_1}}{d\mu_{\text{prior}}} \propto \pi_{\text{like}}(I_1|m). \quad (\text{B.5})$$

Here $\frac{d\mu_{\text{post}}^{I_1}}{d\mu_{\text{prior}}}$ denotes the Radon–Nikodym derivative of $\mu_{\text{post}}^{I_1}$ with respect to μ_{prior} . Assuming that the map $m \mapsto u(m)$ is μ -a.e. well-defined, locally Lipschitz continuous and sufficiently bounded the Bayesian inverse problem is well-posed [29, Corollary 4.4]. The expected information gain (EIG) is then defined as the expectation of the Kullback-Leibler divergence (\mathcal{D}_{KL}) over instances of data I_1 given the noise distribution:

$$\Psi = \mathbb{E}_{I_1} [\mathcal{D}_{KL}(\mu_{\text{post}}(\cdot|I_1) || \mu_{\text{prior}})] \quad (\text{B.6})$$

$$\mathcal{D}_{KL}(\mu_{\text{post}}(m|I_1) || \mu_{\text{prior}}(m)) = \int_{\mathcal{M}} \ln \left(\frac{d\mu_{\text{post}}^{I_1}}{d\mu_{\text{pr}}} \right) \mu_{\text{post}}(dm). \quad (\text{B.7})$$

As was discussed in section 3.5.4, in the case that the forward map $m \mapsto u(m)$ is linear, EIG admits a closed form expression related to the following generalized eigenvalue problem, which is computed at the maximum a posteriori (MAP) point m^* ,

$$m^* = \operatorname{argmax}_{m \in \mathcal{M}} \mu_{\text{post}}(m|I_1) \quad (\text{B.8a})$$

$$\langle D_{mm} \Phi(u(m^*, t); I_0, I_1) \psi_i, \tilde{m} \rangle_{\mathcal{M}} = \lambda_i \langle \mathcal{C}^{-1} \psi_i, \tilde{m} \rangle_{\mathcal{M}} \quad \forall \tilde{m} \in \mathcal{M} \quad (\text{B.8b})$$

$$\Psi = \sum_{i=1}^{\infty} \log(1 + \lambda_i). \quad (\text{B.8c})$$

As noted before m^* coincides with the deterministic inverse problem. While this closed-form expression for EIG is only exact for linear inverse problems, it can be utilized as an approximation for nonlinear problems.

Appendix B.1. Notes on the prior and the function space limit

In the preceding section we focused on quadratic-norm induced Gaussian measures, so we thereby restrict our attention to the \mathcal{R}_{L^2} and \mathcal{R}_{H^1} regularization utilized in the ∞ -IDIC for simplicity. We note that there has been work on formulating infinite-dimensionally consistent TV priors, and the discussion presented here can be naturally extended to that setting [71, 72].

The $L^2(\Omega)$ Tikhonov regularization fails to induce an infinite-dimensionally consistent prior as it is not trace class, since it is induced by the operator $I_{\mathcal{M}}$

$$\mathrm{tr}(I_{\mathcal{M}}) = \sum_{i=1}^{\infty} \langle \psi_i, \psi_i \rangle_{\mathcal{M}} = \sum_{i=1}^{\infty} 1 \not\leq \infty. \quad (\text{B.9})$$

The $H^1(\Omega)$ regularization is associated with the operator $\mathcal{C}^{-1} = -\gamma\Delta + \delta\mathcal{I}_{\mathcal{M}}$ (with homogeneous Neumann boundary conditions), and is only trace class when $\Omega \subset \mathbb{R}^d$ for $d = 1$. Therefore, the Bayesian extension of the \mathcal{R}_{H^1} problem that we introduced, does not exist in the function space limit. It does however, still lead to a meaningful Bayesian extension of an ∞ -IDIC problem *with a fixed mesh representation* as in the numerical results that we discuss in section 3.5.4.

In general one can consider norms based on the $H^k(\Omega)$ semi-norms that do induce meaningful infinite-dimensional Bayesian inverse problems. These are related to the Matérn covariances that are often used in infinite-dimensional Bayesian inverse problems. Particularly, we used them when generating the speckle pattern (see section Appendix A) in addition to generating true modulus fields (i.e., the Gaussian random field in Figure 3). The differential operator

$$\mathcal{A} = \delta I_{\mathcal{M}} - \gamma \nabla \cdot (\Theta \nabla), \quad (\text{B.10})$$

can induce a trace class covariance $\mathcal{C} = \mathcal{A}^{-\alpha}$ (with appropriate boundary conditions), when $\alpha > d/2$ [73]. In (B.10) $\delta, \gamma \in \mathbb{R}$ control the marginal variance and correlation length in the random fields drawn from the associated Gaussian measure, and $\Theta \in \mathbb{R}^{d \times d}$ induces anisotropy in the fields [74]. Consequently, we can take $\delta = 0$, and $\Theta = I$ to recover the $H^k(\Omega)$ semi-norms with $\alpha = k$ and see that we require $k > 1$ for $\Omega \subset \mathbb{R}^2$, but that $k = 2$ works for all $d > 1$. This enforces even more smoothness in the solutions that are favored by the inverse problems (either Bayesian or the analogous deterministic problem). We do not consider this case in numerical experiments, but note that it leads to a meaningful infinite-dimensional Bayesian analogue of the corresponding deterministic ∞ -IDIC problem.