



Optimal experimental design for reliable learning of history-dependent constitutive laws

Kaushik Bhattacharya , Lianghao Cao *, Andrew Stuart 

California Institute of Technology, 1200 E. California Blvd., Pasadena, CA 91125, USA

ARTICLE INFO

Keywords:

Optimal experimental design
Bayesian inverse problems
Machine learning
Solid mechanics
Material testing

ABSTRACT

History-dependent constitutive models serve as macroscopic closures for the aggregated effects of micromechanics. Their parameters are typically learned from experimental data. With a limited experimental budget, eliciting the full range of responses needed to characterize the constitutive relation can be difficult. As a result, the data can be well explained by a range of parameter choices, leading to parameter estimates that are uncertain or unreliable. To address this issue, we propose a Bayesian optimal experimental design framework to quantify, interpret, and maximize the utility of experimental designs for reliable learning of history-dependent constitutive models. In this framework, the design utility is defined as the expected reduction in parametric uncertainty or the expected information gain. This enables *in silico* design optimization using simulated data and reduces the cost of physical experiments for reliable parameter identification.

We introduce two approximations that make this framework practical for advanced material testing with expensive forward models and high-dimensional data: (i) a Gaussian approximation of the expected information gain, and (ii) a surrogate approximation of the Fisher information matrix. The former enables efficient design optimization and interpretation, while the latter extends this approach to batched design optimization by amortizing the cost of repeated utility evaluations. Our numerical studies of uniaxial tests for viscoelastic solids show that optimized specimen geometries and loading paths yield image and force data that significantly improve parameter identifiability relative to random designs, especially for parameters associated with memory effects.

1. Introduction

1.1. Background and motivation

Constitutive relations close the balance equations of continuum mechanics by mapping the kinematic histories of continua to their internal response. To determine the constitutive relations for a system of interest, one first selects a constitutive model with unspecified parameters based on physical considerations. Then, the system-specific parameter values are identified from experiments conducted under controlled and simplified scenarios. The calibrated constitutive models are then deployed to predict quantities of interest in more complex, realistic scenarios, informing critical decision-making for the system. The quality of these predictions and the subsequent decisions hinges on the quality of the parameter identification. Consequently, it is of great interest to develop techniques that reliably identify constitutive parameters from experiments.

* Corresponding author.

E-mail addresses: bhatta@caltech.edu (K. Bhattacharya), lianghao@caltech.edu (L. Cao), astuart@caltech.edu (A. Stuart).

<https://doi.org/10.1016/j.cma.2026.119022>

Received 12 March 2026; Received in revised form 17 April 2026; Accepted 20 April 2026

Available online 25 April 2026

0045-7825/© 2026 Elsevier B.V. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

The parameters defining constitutive relations are often not directly measurable. Instead, the parameters are identified by solving the inverse problem: finding parameter values that minimize the discrepancy between experimental data and model predictions. These predictions are generated by a forward model that consists of two parts: (i) the balance equations, adapted to the specific experimental setup, such as geometry and boundary conditions; and (ii) an observation model that predicts the experimental data from the solutions of the balance equations. Three major challenges arise when solving this inverse problem in practice:

- (i) *Computational Complexity*. Repeated solution of the balance equations, required to solve the inverse problem, can render the computations intractable under time and computing resource constraints.
- (ii) *Robustness*. There are inevitable discrepancies between the forward model and the reality of experiments, such as data noise and model misspecification. If the inverse problem is unstable, these discrepancies can be amplified, leading to inaccurate or unphysical parameter estimates.
- (iii) *Data Quality*. Experimental data is often limited and noisy and, as a consequence, may not contain sufficient information to constrain the parameters. This leads to uncertain or unreliable parameter identification, thereby limiting the predictive power of the calibrated constitutive model.

These three challenges interact. Lack of robustness amplifies the downstream effects of uncertain or unreliable parameter estimates [1,2]. This can be addressed by the Bayesian approach to inversion [3], which quantifies uncertainty and improves robustness, but it places a heavy computational burden. Finally, we note that experimental design [4,5] can mitigate issues with data quality but adds to the computational burden.

In the context of solid-like materials, various experimental and computational techniques have been developed to address the challenges above; however, data-quality limitations remain largely unresolved. For example, conventional material testing, such as uniaxial loading and shear rheometry, prioritizes simplicity and robustness; however, data quality is often limited due to highly constrained experimental setups [6]. Emerging techniques that use full-field observations [7] capture detailed, heterogeneous displacement fields from images, thereby increasing the information capacity of the data for parameter identification. However, exploiting this capacity to improve parameter identifiability remains challenging, as it increases the complexity of the experimental design—see Material Testing 2.0 [6] and a review of related work on material testing for parameter identification in Section 1.3.

This paper addresses data-quality limitations in constitutive parameter identification using optimal experimental design. A specific focus is on history-dependent constitutive models in which stress depends on deformation history through a memory representation, such as hereditary integrals, fractional time derivatives, or internal state variables. In inelastic materials, the history dependence reflects the combined effects of dissipative microscale mechanisms, such as dislocation glide, chain segment relaxation, and microvoid growth. Identifying their constitutive parameters requires experimental data characterizing the material response under a diverse set of deformation histories. Designing a small number of experiments that generate such information-rich data is challenging for two primary reasons: (i) it is often unclear what type of deformation histories are most relevant to parameter identifiability; and (ii) most experiments cannot directly impose heterogeneous deformation histories pointwise to enrich the resulting full-field observations. For certain constitutive models, we have insights into experimental designs that can isolate or amplify signals of operative mechanisms in experimental data, such as stress relaxation in viscoelastic solids. However, gaining such insight is not straightforward for general, nonlinear constitutive models, such as neural network models [8–10], which have high expressivity but low interpretability. This motivates our focus and the contributions we now overview.

1.2. Main contributions and outline of paper

We introduce a Bayesian framework to quantify, interpret, and maximize the utility of experimental designs for reliable parameter identification. Departing from intuition-driven or heuristic strategies, we employ the expected information gain to directly quantify parameter identifiability for a given design. Furthermore, we address the limitations of data-driven [11] and sensitivity-based approaches [12,13], which typically require exploratory physical testing with suboptimal designs to inform the optimization process. In contrast, our framework explores the joint design-parameter space using simulated experimental outcomes, thereby eliminating the reliance on exploratory testing and reducing total testing time and costs. Our specific contributions are as follows:

- (CI) We formulate a Bayesian optimal experimental design framework to improve the parameter identifiability of history-dependent constitutive models from experimental data. It uses expected information gain to assess data quality and to enable *in silico* design optimization.
- (CII) We employ a Bayesian D-optimal design utility based on a Gaussian approximation of the expected information gain for efficient design optimization and interpretation.
- (CIII) We propose and numerically validate a method that amortizes the cost of batched design optimization by surrogate modeling of the Fisher information matrix.
- (CIV) We study the proposed framework for designing uniaxial tests for viscoelastic solids. The results show that the optimized specimen shape and loading path yield high-quality image and force data, resulting in more reliable parameter estimates than those obtained from randomly selected tests. The implementation is available through the following link: https://github.com/lcao11/constitutive_oed/.

The paper is organized as follows. We conclude this section with a review of related work in Section 1.3. Section 2 introduces history-dependent constitutive laws, particularly with memory representation based on internal state variables. Section 3 presents the Bayesian formulation of constitutive parameter identification from experimental data. Section 4 introduces Bayesian optimal experimental design, where the EIG design utility (CI) and its approximations (CII and CIII) are derived. Sections 5 and 6 present our numerical study on the design of uniaxial tests to learn linear and nonlinear viscoelasticity (CIV), respectively. Section 7 contains concluding remarks. The appendices contain details of the numerical implementation.

1.3. Related work

We provide further details concerning each of the contributions in this paper and review related work. The purpose of this review is to set our contributions in the context of the literature rather than to serve as an exhaustive survey. We first address Bayesian optimal experimental design for material testing (CI) in Section 1.3.1 and Bayesian D-optimal design (CII) in Section 1.3.2. We then discuss surrogate modeling for experimental design (CIII) in Section 1.3.3, followed by advanced material testing for parameter identification (CIV) in Section 1.3.4.

1.3.1. Bayesian optimal experimental design for material testing

Details of (CI). We use the expected information gain (EIG) as a baseline measure of the utility of experimental design for reliable parameter identification. The EIG is based on a Bayesian formulation of parameter identification that yields a posterior distribution rather than a point estimate from the experimental data. The EIG quantifies the expected reduction in parametric uncertainty, measured by information entropy, from model-predicted data at the candidate design. It can also be interpreted as the mutual information between the parameters and data. \diamond

Expected information gain. Lindley [14] first proposed using the EIG to measure experiment utility, a concept now standard in Bayesian experimental design; see reviews on this topic in [4,5,15,16]. For material testing, the EIG has been shown to be effective for designing loading steps for biaxial testing [17], Weber number and bubble radius for inertial microcavitation rheometry [18], and grain orientation for spherical indentation [19]. These works employ the incremental EIG for myopic sequential design, in which the posterior from the previous experiment serves as the prior for generating the next design. They use simplified but approximate forward models that circumvent the balance equations. In contrast, our forward models incorporate them to fully assimilate image data that capture spatially heterogeneous material responses. Additionally, our focus is on non-adaptive batched design, which can serve as an incremental step in sequential design. Finally, the batched and myopic sequential designs are special cases of sequential designs formulated as Markov decision processes [20,21]. \diamond

Stress state entropy. Recently, information-theoretic metrics have been proposed to design specimen geometries for uniaxial and biaxial loading that maximize the heterogeneity of induced stress fields. Prominently, the stress state entropy framework [22,23] employs information entropy to quantify the diversity of stress invariants, such as triaxiality and Lode angle, within a single experiment. This framework assumes that maximizing stress state diversity is equivalent to maximizing the information content for parameter identification. In contrast, we consider the EIG, which directly assesses parameter identifiability and quantifies the mutual information between the data and the parameters. \diamond

1.3.2. Bayesian D-optimal design based on Gaussian posterior approximation

Details of (CII). The EIG is expensive to estimate and can be difficult to interpret. We therefore adopt the Bayesian D-optimal design utility derived from the EIG using a Gaussian posterior approximation. This replaces the nested loop of forward model evaluations for EIG estimation with a single loop of Fisher information matrix (FIM) evaluations, thereby substantially reducing the cost of utility evaluations. Moreover, because Gaussian marginals are available in closed form, the utility can be evaluated for any parameter subset at negligible additional cost, enabling quantitative diagnosis of how design choices affect recovery of parameters governing specific behaviors such as anisotropy and stress relaxation. \diamond

Sensitivity-based design. Our approach is connected to sensitivity-based optimal design of specimen geometry [24,25] and loading path [12] for uniaxial or biaxial testing. Their design formulations are equivalent to classical D-optimal designs [5,13], which maximize the logarithm of the determinant of the FIM for a linear or linearized forward model; this connection was recently noted by Fayad et al. [12]. For nonlinear models, design optimization involves an iterative process of experimentation, parameter identification, and optimization using the linearized model at the parameter point estimate. This procedure faces a dilemma: reliable estimates are required for good design, yet good design requires reliable estimates. In contrast, our approach does not require a point estimate of the parameter to initiate the optimization. We formulate a risk-neutral optimization under prior uncertainty, eliminating the need for exploratory physical testing with suboptimal designs. \diamond

Gaussian posterior approximation. There are two commonly-used approximations derived from a quadratic expansion of the log-posterior centered at different parameter estimates: (i) the Laplace approximation [26–30], centered at the maximum a posteriori (MAP) estimator, and (ii) a data-independent local approximation [27,29,31], centered at the data-generating parameters. Although (i) typically provides higher fidelity, it requires parameter identification for each data realization. By contrast, (ii) is cheaper because it avoids this optimization procedure and exploits the known data-generating parameters. When the forward model evaluation requires solving the balance equations, the repeated parameter identification for each utility evaluation creates a computational bottleneck in design optimization. This issue is exacerbated when history-dependent constitutive models are used. This motivates our use of (ii). \diamond

Bayesian D-optimal design. The connection between Bayesian D-optimal design utility and the EIG under a Gaussian posterior approximation is well established [4]. Such approximation is widely used to mitigate the prohibitive cost of nested Monte Carlo estimators of EIG [32,33]. In practice, the Gaussian approximation is employed either as (i) an analytic replacement of the posterior density in EIG approximation [27,29,30,34] or (ii) an importance sampling proposal to accelerate EIG estimation [26,28,31]. The Bayesian D-optimal utility emerges in case (i), which is adopted in this work. \diamond

1.3.3. Surrogate modeling for optimal experimental design

Details of (CIII). Material testing often involves repeated experiments, where a batch of experiments with random designs may yield more useful information than a single optimized experiment. This motivates batched experimental design optimization. However, its computational cost increases linearly with the batch size. To amortize this cost, we train a surrogate model of the FIM, incurring a one-time cost to generate training samples that is independent of the batch size. Then the trained surrogate model predicts the FIMs during batched design optimization, making it tractable even with large batch sizes. \diamond

Surrogate utility and surrogate forward model. Most surrogate modeling approaches for experimental design either learn the utility function [35,36] or the forward model [37–40]. Surrogate utility emulates the design utility at a fixed batch size, and its training cost depends on that batch size. In our numerical studies, we also use a Gaussian process to emulate the utility for design optimization (i.e., Bayesian optimization). In contrast, the training costs of surrogate forward models are independent of batch size, and the surrogate can be used to evaluate likelihood across multiple experiments. However, training surrogate forward models can be challenging when the dimensions of the parameters, design variables, or data are high. For advanced material testing, snapshots of image data or full-field observations contain millions of data points, making it difficult to construct surrogate forward models without resorting to additional data compression techniques. \diamond

Surrogate Fisher information matrix. The FIM size is independent of the data dimension, yet scales quadratically with the parameter dimension. When the parameter dimension is much smaller than the data dimension, it is easier to construct surrogate FIMs than surrogate forward models. We therefore adopt this approach in our work, training a neural network to emulate the FIM as a function of the parameter and design variables. Similarly, Lan et al. [41] use a Gaussian process to emulate the FIM for fast metric tensor predictions in MCMC. Recent works explore using the derivative-informed neural operator [42] to emulate both the forward model and its Jacobian, which is then used for FIM predictions in geometric MCMC [43] and optimal experimental design [40]. \diamond

1.3.4. Advanced material testing for parameter identification

Details of (CIV). We provide numerical case studies on uniaxial testing of linear and nonlinear viscoelastic solids. The experiment yields reaction force measurements and image snapshots of the deformed specimen as data for parameter identification. The following experimental designs are considered simultaneously:

- (i) *Specimen Geometry.* We consider a thin, flat specimen with a centered elliptical hole. The hole's aspect ratio and orientation are design variables.
- (ii) *Loading Path.* The axial loading path is prescribed by the strain values at discrete time points with interpolation. These strain values are design variables.

We note that these design variables are chosen to demonstrate the applicability of our framework to advanced material testing using image or full-field observations. Our design optimization formulation is general and not limited to these specific settings. \diamond

Full-field observations. Digital image correlation (DIC) [44–46] infers displacement fields by correlating images of a speckled surface before and after deformation. It has also been extended to obtain out-of-plane [47,48] and 3D displacement fields [49–51]. Full-field observations have high information capacity for parameter identification. For example, one can probe many different deformation histories in a single test using specimens with complex geometries [24]. This information capacity can be further leveraged through topology optimization of specimen geometry [25,52,53]. In our numerical case studies, we also incorporate specimen geometry into the experimental design. \diamond

Parameter identification from full-field observations. Numerical methods such as the finite element model updating (FEMU) method [54, 55] and the virtual fields (VFM) method [56,57] can exploit full-field observations for parameter identification; see [7] for a review. Other approaches that treat parameter identification as PDE-constrained inverse problems [58,59] are most relevant to this work. These inverse-problem formulations employ weighted least-squares data misfits, which are equivalent to MAP estimation problems under a Bayesian formulation with additive Gaussian noise [2,3]. \diamond

Parameter identification from images. In our numerical studies, images are used directly for parameter identification rather than the displacements inferred from them. This aligns with integrated DIC and other image-based approaches [60–63], which avoid inverting for displacement fields and directly minimize image misfit or the grayscale residual. This integrated approach has been shown to yield more accurate and reliable parameter estimates than the two-step approach of first applying DIC, then performing parameter identification [64]. The motivation for our use of image data is the simplicity of image noise models, such as additive and uncorrelated Gaussian or Poisson–Gaussian noise [65,66], which leads to closed-form likelihoods and allows for FIM evaluations using the Jacobian of the forward model [67]. In contrast, the noise models for the inferred displacement fields are nonphysical and nonlinear [68] with spatial correlation [69]. \diamond

Bayesian formulation of parameter identification. The theory and computational methods for Bayesian inverse problems have seen substantial development in the past two decades; see, e.g., [1–3]. It has also been applied to parameter identification using conventional testing; see, e.g., [70]. Recent works have also applied this formulation to full-field observations [71–74] or image data [75]. The Bayesian formulation considered in our work is general and not restricted to full-field observations; for instance, our numerical case studies consider both force and image data. Our formulation involving image data is related to Gaynutdinova et al. [75], which extends integrated DIC to define a pixel-level likelihood. \diamond

2. History-dependent constitutive laws

In this section, we present history-dependent constitutive laws based on the internal variable theory [76] to model the isothermal behavior of materials. In this framework, the state of a material is described not only by the deformation gradient F , but also by a set of scalar or tensorial internal states, $\{\alpha_i\}_{i=1}^N$ (denoted $\alpha = \{\alpha_i\}$ for brevity, and similarly for other variables.) The evolution of these internal states is typically needed to account for changes in the material microstructure. However, we emphasize that the experimental design methodology introduced in this work is not restricted to internal variable theory and can also be applied to other memory representations, such as time integrals and fractional derivatives.

In Section 2.1, we describe some common modeling considerations arising in history-dependent constitutive laws. Sections 2.2 and 2.3 then describe, respectively, the specific cases of linear viscoelasticity and anisotropic finite-strain viscoelasticity used in our numerical examples. For a general discussion of viscoelasticity, we refer the reader to [77,78].

2.1. Common modeling considerations

Helmholtz free energy. The Helmholtz free energy Ψ , a measure of stored material energy, is assumed to depend only on F and α :

$$\Psi = \Psi(F, \{\alpha_i\}).$$

The first Piola–Kirchhoff (PK1) stress, denoted as P , is the stress measure that is work-conjugate to F . It is often assumed to be additively decomposed into an equilibrium component, defined through the free energy, and a non-equilibrium component,

$$P = P^{\text{eq}} + P^{\text{neq}}, \quad P^{\text{eq}} = \partial_F \Psi(F, \{\alpha_i\}).$$

\diamond

Dissipation inequality and kinetic laws. For an isothermal process, the second law of thermodynamics requires that the rate of internal dissipation be non-negative:

$$P^{\text{neq}} : \dot{F} + \sum_{i=1}^{n_\alpha} A_i \dot{\alpha}_i \geq 0, \quad A_i := -\partial_{\alpha_i} \Psi(F, \{\alpha_j\}). \quad (1)$$

where \dot{F} denotes the rate of change of the deformation gradient, and A_i is the thermodynamic driving force conjugate to the internal state α_i defined as the partial derivative of free energy.

The dissipation inequality constrains the admissible forms for the dissipative stress P^{neq} and the kinetic laws that govern the evolution of the internal states. We consider the following commonly adopted form of kinetic law, defined through function $\mathcal{G} = \{\mathcal{G}_i\}$:

$$\dot{\alpha}_i = \mathcal{G}_i(F, \dot{F}, \{\alpha_j\}, \{A_j\}).$$

\diamond

Hereditary operator. Together, the modeling choices for the free energy potential Ψ , kinetic laws $\{\mathcal{G}_i\}$, and the initial condition of the internal states form a constitutive model. This model can be viewed as a hereditary operator \mathcal{M} that predicts the stress state in the reference frame from the history of deformation:

$$P(t) = \mathcal{M}(\{F(s)\}_{s \in [0,t]}, \theta),$$

where θ represents the constitutive parameters. \diamond

Material frame indifference. The Helmholtz free energy Ψ and the kinetic laws \mathcal{G}_i should satisfy the material frame indifference principle, which can be achieved by defining Ψ and $\{\mathcal{G}_i\}$ through objective tensors that depend on F and \dot{F} , such as the right Cauchy–Green tensor $C := F^T F$, and the rate of deformation tensor $D = \text{sym}(\dot{F} F^{-1})$, where $\text{sym}(A) = (A + A^T)/2$. \diamond

Dissipative potential. A common approach for designing the kinetic laws to meet the constraint of the dissipation inequality is to postulate a dissipation potential Φ defined as

$$\Phi = \Phi(F, \dot{F}, \{\alpha_i\}, \{\dot{\alpha}_i\}),$$

from which the dissipative quantities are then derived:

$$P^{\text{neq}} = \partial_{\dot{F}} \Phi(F, \dot{F}, \{\dot{\alpha}_i\}, \{\alpha_i\}), \quad A_i = \partial_{\dot{\alpha}_i} \Phi(F, \dot{F}, \{\alpha_i\}, \{\dot{\alpha}_j\}). \quad (2)$$

The convexity of Φ with respect to the rate variables \dot{F} and $\{\dot{\alpha}\}$ ensures the satisfaction of dissipation inequality. The kinetic law for each internal state is then found by equating the energetic driving force in Eq. (1) with the dissipative force in Eq. (2). \diamond

2.2. Example I: Linear viscoelasticity

We consider three-dimensional linear viscoelasticity. Let $\alpha_i : [0, T] \rightarrow \mathbb{R}_{\text{sym}}^{3 \times 3}$ for $1 \leq i \leq N$ denote a set of tensorial internal variables. Let $\varepsilon = \frac{1}{2}(F + F^T) - I$ denote the infinitesimal strain tensor. The Helmholtz free energy and the dissipative potential are given by

$$\Psi(\varepsilon, \alpha; \{\mathbb{C}_i\}) = \frac{1}{2} \varepsilon : \mathbb{C}_0 : \varepsilon + \frac{1}{2} \sum_{i=1}^N (\varepsilon - \alpha_i) : \mathbb{C}_i : (\varepsilon - \alpha_i),$$

$$\Phi(\dot{\alpha}; \{\tau_i\}, \{\mathbb{C}_i\}) = \frac{1}{2} \sum_{i=1}^N \tau_i \dot{\alpha}_i : \mathbb{C}_i : \dot{\alpha}_i,$$

where $\{\mathbb{C}_i\}$ are symmetric and coercive 4th-order tensors representing the elastic and viscous effects and $\{\tau_i\}$ are the relaxation times. The resulting constitutive model is the generalized Maxwell model that predicts the stress $\sigma \approx P$ in $t \in [0, T]$:

$$\sigma(t) = \mathbb{C}_0 : \varepsilon(t) + \sum_{i=1}^N \mathbb{C}_i : (\varepsilon(t) - \alpha_i(t)),$$

$$\dot{\alpha}_i(t) = \tau_i^{-1} (\varepsilon(t) - \alpha_i(t)), \quad 1 \leq i \leq N,$$

$$\alpha_i(0) = 0, \quad 1 \leq i \leq N.$$

This constitutive model is equivalent to the following memory kernel model under the Prony series expansion when $\varepsilon(0) = 0$:

$$\sigma(t) = \mathbb{C}_0 : \varepsilon(t) + \int_0^t \mathbb{K}(t-s; \{\mathbb{C}_i\}, \{\tau_i\}) : \dot{\varepsilon}(s) ds,$$

$$\mathbb{K}(t) = \sum_{i=1}^N \mathbb{C}_i \exp(-t/\tau_i).$$

In this linear viscoelastic constitutive model, at most $d_\theta = 22N + 21$ parameters need to be determined, since each 4th-order tensor contains at most 21 independent components. For the numerical example in Section 5, we use this model with $N = 2$ and orthotropic plane stress, leading to 11 parameters.

2.3. Example II: Anisotropic finite-strain viscoelasticity

We consider a three-dimensional finite-strain, anisotropic nonlinear viscoelastic solid with N parallel viscous branches [79]. The equilibrium response is hyperelastic with fiber reinforcement, and the rate dependence arises from non-linear viscous branches driven by a dual dissipative potential. In this model, the internal state variables are $\{F_i^v\}_{i=1}^N$, each represents the viscous deformation gradient for the i th viscous branch. A multiplicative decomposition of the deformation gradient is assumed for each branch: $F = F_i^e F_i^v$, where F_i^e is the elastic deformation gradient for the i th branch. The total stress response is the sum of an equilibrium response and N non-equilibrium responses. The Helmholtz free energy is given by an equilibrium energy Ψ^{eq} and a set of non-equilibrium energies $\{\Psi_i^{\text{neq}}\}$:

$$\Psi(F, \{F_i^v\}) = \Psi^{\text{eq}}(C) + \sum_{i=1}^N \Psi_i^{\text{neq}}(C_i^e),$$

where $C_i^e = (F_i^e)^T F_i^e$ is the strain tensor of the intermediate configuration for each viscous branch.

Let $(\cdot) = J^{-2/3}(\cdot)$ be the isochoric operator where $J = \det(F)$. The relevant invariants of \bar{C} are $\bar{I}_1 = \text{tr}(\bar{C})$ and $\bar{I}_4 = \bar{C} : A_0$, where $A_0 = a_0 \otimes a_0$ denotes the structure tensor and a_0 is the fiber orientation unit vector. The equilibrium energy is given by

$$\Psi^{\text{eq}}(C; \mu_\infty, \kappa, k_1, k_2, a_0) = \frac{\mu_\infty}{2} (\bar{I}_1 - 3) + \frac{\kappa}{2} (\ln J)^2 + \frac{k_1}{2k_2} \left(\exp(k_2 \langle \bar{I}_4 - 1 \rangle^2) - 1 \right) + \frac{\tau k_1}{2} \langle 1 - \bar{I}_4 \rangle^2,$$

where μ_∞ is the shear modulus of the matrix (i.e., the background material), κ is the bulk modulus, k_1 is the initial (small-strain) fiber stiffness, k_2 is the stiffening rate, and $\langle \cdot \rangle$ is the Macaulay bracket defined as $\langle x \rangle := \max\{0, x\}$. The fiber energy includes both tensile and compressive components [80]. The exponential term is active only when fibers are in tension ($\bar{I}_4 > 1$), while a small quadratic compressive penalty, scaled by $\tau = 0.1$, is active in compression ($\bar{I}_4 < 1$) to prevent numerical instabilities (arising, for example, from buckling). To capture the reorientation of fibers during viscous flow, the non-equilibrium fiber invariants are defined in the intermediate configuration:

$$\bar{I}_{4,i}^e = \bar{C}_i^e : A_i^v, \quad \text{with} \quad A_i^v = a_i^v \otimes a_i^v, \quad a_i^v = \frac{F_i^v a_0}{\|F_i^v a_0\|}.$$

Physically, a_i^v represents the fiber direction convected by the viscous flow F_i^v prior to elastic stretching [81]. We define $\bar{I}_{1,i}^e = \text{tr}(\bar{C}_i^e)$ and assume the viscous flow is isochoric, which implies $J_i^v = 1$ and $\det(C_i^e) = J^2$. The forms of the non-equilibrium energies are:

$$\Psi_i^{\text{neq}}(C_i^e; \mu_i, k_{1,i}, k_{2,i}, a_0) = \frac{\mu_i}{2} (\bar{I}_{1,i}^e - 3) + \frac{k_{1,i}}{2k_{2,i}} \left(\exp(k_{2,i} \langle \bar{I}_{4,i}^e - 1 \rangle^2) - 1 \right) + \frac{\tau k_{1,i}}{2} \langle 1 - \bar{I}_{4,i}^e \rangle^2,$$

where $\{\mu_i\}$, $\{k_{1,i}\}$, $\{k_{2,i}\}$ are parameters of the non-equilibrium energies.

The dissipation for each viscous branch is governed by the conjugate pair (M_i^v, L_i^v) , where $L_i^v = \dot{F}_i^v(F_i^v)^{-1}$ is the viscous velocity gradient and $M_i^v = C_i^e S_i^e$ is the Mandel stress for the i th branch, with $S_i^e = 2\nabla\Psi_i^{\text{neq}}(C_i^e)$ being the thermodynamic driving force. The kinetic law is prescribed by dual dissipative potentials $\{\Phi_i^*\}$ [82,83]:

$$\Phi_i^*(M_i^v; \gamma_i, m_i) = \frac{\gamma_i}{m_i + 1} \|\text{dev}(\text{sym}(M_i^v))\|^{m_i+1},$$

where $\{\gamma_i\}$ are the rate coefficients, $\{m_i\}$ are the rate-sensitivity exponents.

The explicit form of the constitutive law for $t \in [0, T]$ is given by

$$\begin{aligned} P(t) &= F(t)S(t), \\ S(t) &= 2\nabla\Psi^{\text{eq}}(C(t)) + 2 \sum_{i=1}^N F_i^v(t)^{-1} \nabla\Psi_i^{\text{neq}}(C_i^e(t)) F_i^v(t)^{-\top}, \\ M_i^v(t) &= 2C_i^e(t) \nabla\Psi_i^{\text{neq}}(C_i^e(t)), & 1 \leq i \leq N, \\ L_i^v(t) &= \gamma_i \|\text{dev}(\text{sym}(M_i^v))\|^{m_i-1} \text{dev}(\text{sym}(M_i^v)), & 1 \leq i \leq N, \\ \dot{F}_i^v(t) &= L_i^v(t) F_i^v(t), & 1 \leq i \leq N, \\ F_i^v(0) &= I, & 1 \leq i \leq N. \end{aligned}$$

In this nonlinear viscoelastic model, there are in total $5N + 4$ constitutive parameters to be determined from experimental data.

3. Bayesian inverse problems

In this section, we formulate Bayesian inverse problems (BIPs) for identifying constitutive parameters from experimental data. We first define the forward model, comprising the balance equation (Section 3.1) and observation operator (Section 3.2). We then present the model-constrained BIP framework in Section 3.3 and extend it to multiple experiments in Section 3.4.

3.1. The balance equations

We consider an open, bounded domain $\Omega_0 \subset \mathbb{R}^3$ occupied by a material undergoing deformation. We assume the deformation is described by a displacement field $u(X, t)$, where $X \in \Omega_0$ denotes and $t \in [0, T]$ denotes time with a terminal time $T > 0$. We are interested in the material response to external stimuli, such as body forces, boundary traction forces, boundary conditions, or initial conditions. Let $z \in \mathbb{R}^{d_z}$ denote the experimental design variables. The governing equation for the displacement field is given by

$$\rho \ddot{u}(X, t) - \text{Div}_X(P(X, t)) = f(X, t; z), \quad (X, t) \in \Omega_0(z) \times [0, T], \tag{3a}$$

$$P(X, t) = \mathcal{M}(\{F(X, s)\}_{s \in [0, t]}; \theta), \quad (X, t) \in \Omega_0(z) \times [0, T], \tag{3b}$$

$$P(X, t)n_0(X) = \tau_n(X, t; z), \quad (X, t) \in \Gamma_N(z) \times [0, T], \tag{3c}$$

$$u(X, t) = u_D(X, t; z), \quad (X, t) \in \Gamma_D(z) \times [0, T], \tag{3d}$$

$$u(X, 0) = u_0(X; z), \quad X \in \Omega_0(z), \tag{3e}$$

$$\dot{u}(X, 0) = v_0(X; z), \quad X \in \Omega_0(z), \tag{3f}$$

where ρ is the material density, f is the external force, n_0 is the outward normal vector of the material domain, τ_n is the boundary traction defined over the material domain boundary $\Gamma_N \subseteq \partial\Omega_0$, u_D is the Dirichlet boundary defined over the material domain boundary $\Gamma_D \subseteq \partial\Omega_0$, and u_0 and v_0 are the initial condition for the displacement field and its velocity. In abstract form, the solution operator of the governing equation is denoted as the map from the experimental design variables and the constitutive parameters to the trajectory of the material displacement:

$$u = S(\theta, z). \tag{4}$$

Example: Uniaxial tests with loading path and specimen shape design. We consider a uniaxial test of a flat, thin specimen under the 2D plane-stress assumption. Let $\Omega_0(z_s) = ([0, L_1] \times [0, L_2]) \setminus \Omega_H(z_s)$ denote a rectangular-shaped material surface with holes $\Omega_H(z_s) \subset [0, L_1] \times [0, L_2]$ parameterized by the shape parameters $z_s \in \mathbb{R}^{d_{z_s}}$. The material is clamped at the left boundary (zero displacement) and pulled on the right boundary according to a loading path $u_l : [0, 1] \rightarrow [0, u_{\max}]$. We take $u_l(0) = 0$ and $u_{\max} > 0$ as the maximum loading position. Let $z_l \in [0, u_{\max}]^{d_{z_l}}$ denote the equally spaced control points of the loading path. We prescribe the loading path via interpolation, i.e.,

$$u_l(t_i; z_l) = [z_l]_i, \quad t_i = i/d_{z_l}, \quad i = 1, \dots, d_{z_l}.$$

Under the quasi-static assumption, this setup is obtained from the balance equations above by setting $\rho \ddot{u} = 0$, $f = 0$, and $\tau_n = 0$, while retaining the same hereditary constitutive law. The resulting equilibrium equation is

$$-\text{Div}_X(P(X, t)) = 0, \quad (X, t) \in \Omega_0(z_s) \times [0, T],$$

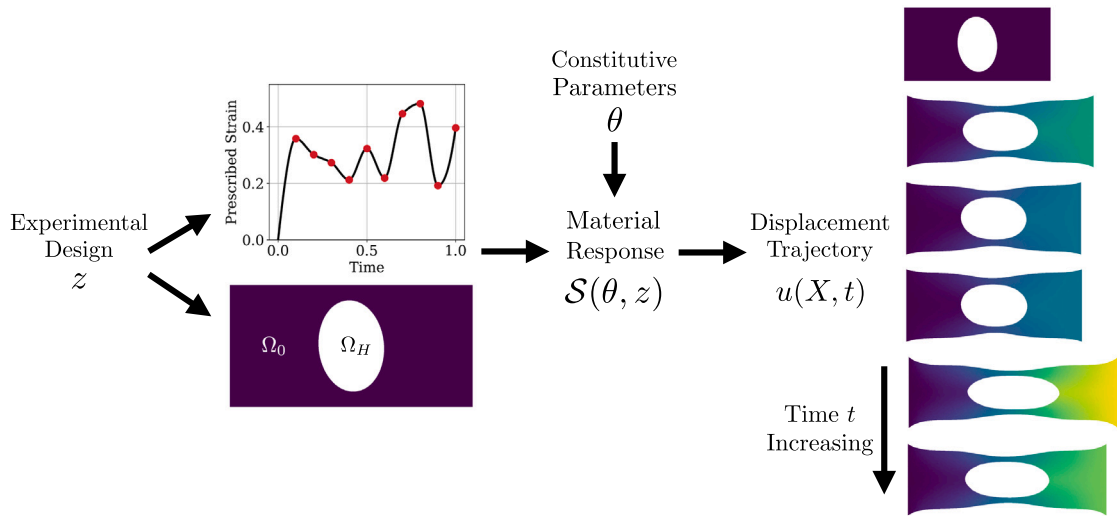


Fig. 1. A visualization of a model for the described uniaxial test with loading path and specimen shape design. The dots on the loading path, i.e., the prescribed strain, are the control points. The color on the deformed material represents the magnitude of the displacement.

with Dirichlet boundary conditions

$$u(X, t) = 0 \quad \text{on } \{0\} \times [0, L_2], \quad u(X, t) = (u_l(t; z_l), 0) \quad \text{on } \{L_1\} \times [0, L_2],$$

a traction-free boundary condition on the remaining boundary, and zero initial displacement. Fig. 1 shows examples for the input and output of the solution operator $u = S(\theta, z)$ for this setup of uniaxial testing. \diamond

3.2. The observation operator and the forward model

During experiments, observations of the material’s response are collected, including images and reaction forces. The observations can also be designed, for example, with a speckled pattern, and thus may depend on the design variables z . In an abstract form, we model the experimental data $y \in \mathbb{R}^{d_y}$ as

$$y = \mathcal{O}(u, \theta, z, \eta), \quad \eta \sim \pi_\eta. \tag{5}$$

Here \mathcal{O} is the observation operator that returns the experimental data given the displacement fields, the constitutive parameters, and the experimental design. The unknown noise $\eta \in \mathbb{R}^{d_y}$ follows the noise distribution π_η .

By combining the solution operator of the governing equation, encapsulated in (4), with the observation model (5), we arrive at our forward model F :

$$y = F(\theta, z, \eta), \quad F(\theta, z, \eta) := \mathcal{O}(S(\theta, z), \theta, z, \eta). \tag{6}$$

Example: Image data. We consider image data that captures a 2D surface of a deformed material. The surface is painted with a speckled pattern, and this data is often used in inferring displacement fields via DIC; see Section 1.3. Here, we define a model for a snapshot of image data produced at time $t \in [0, T]$, with a large field of view that covers the full range of deformation in the experiment. The details of this model are provided in Appendix A.

Let $I_0 \in [0, 1]^{n_W \times n_H}$ denote a reference image of the speckled pattern painted on a flat material surface $\partial\Omega_I \subset \Omega$. Pixels located outside the surface (i.e., the environment) are masked to zero intensity. First, we define an image deformation operator \mathcal{W} that uses the displacement value at each pixel of the reference image and a push forward method to obtain an image of the deformed material $I_{\text{pred}} \in [0, 1]^{n_W \times n_H}$ with masked environment, i.e.,

$$I_{\text{pred}} = \mathcal{W}(u(\cdot, t)|_{\partial\Omega_I}, I_0) \in [0, 1]^{n_W \times n_H}, \tag{7}$$

where $u(\cdot, t)|_{\partial\Omega_I}$ denotes the displacement field on the surface at time t . The deformed image is perturbed with white noise that has a variance σ_η^2 , and a soft mask $\mathcal{A}(I_{\text{pred}}) \in [0, 1]^{n_W \times n_H}$ is then applied to remove the intensity of pixels corresponding to the environment:

$$y = \mathcal{A}(I_{\text{pred}}) \odot (I_{\text{pred}} + \sigma_\eta \eta), \quad [\eta]_{ij} \sim \mathcal{N}(0, 1),$$

where \odot denotes pixel-wise multiplication. Fig. 2 contains examples of the observation operator $y = \mathcal{O}(u, I_0, t, \eta)$ defined through the procedure above. Finally, we note that one may include I_0 and t as design variables. In our numerical examples, however, we choose to fix them and design the specimen geometry and loading path. \diamond

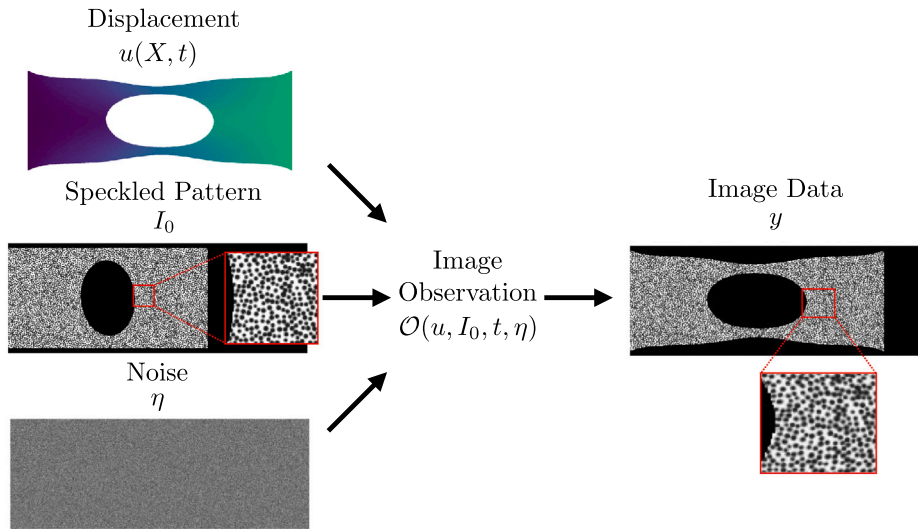


Fig. 2. A visualization of the observation operator for a single snapshot of image data at time t for the uniaxial testing described in Section 3.1.

3.3. Model-constrained Bayesian inverse problems

We now adopt a Bayesian probabilistic framework for inferring the constitutive parameters from the data. Let Θ denote the parameter random vector following the prior distribution π_{Θ} . We may define a data random vector Y for experimental design variables z as

$$Y(z) = \mathcal{F}(\Theta, z, \eta), \quad \Theta \sim \pi_{\Theta}, \quad \eta \sim \pi_{\eta}.$$

We denote the conditional distribution of $Y(z)$ given θ , i.e., the likelihood, as $\pi_{Y(z)|\theta}(\cdot | \theta)$. A common form of the likelihood can be derived from the additive noise model, which assumes that there exists a parameter-to-observable map \mathcal{F}_{pto} such that $Y(z) = \mathcal{F}_{\text{pto}}(\theta, z) + \eta$. In this case, we have

$$\pi_{Y(z)|\theta}(y | \theta) := \pi_{\eta}(y - \mathcal{F}_{\text{pto}}(\theta, z)).$$

The posterior distribution of the constitutive parameters, denoted by $\pi_{\Theta|Y(z)}$, is defined via the Bayes' rule

$$\pi_{\Theta|Y(z)}(\theta|y) = \frac{\pi_{Y(z)|\theta}(y | \theta) \pi_{\Theta}(\theta)}{\pi_{Y(z)}(y)},$$

where the data marginal distribution or marginal likelihood $\pi_{Y(z)}$ is given by

$$\pi_{Y(z)}(y) = \mathbb{E}_{\theta \sim \pi_{\Theta}} [\pi_{Y(z)|\theta}(y | \theta)]. \tag{8}$$

Given the observed data y produced with design z , the goal of the Bayesian inverse problem is to characterize the posterior $\pi_{\Theta|Y(z)}(\cdot | y)$, which encapsulates the parametric uncertainty after the experiment.

3.4. Bayesian inverse problems for multiple experiments

Now we consider the case where N_E experiments are conducted with the same setup described by the model \mathcal{F} but with a set of different experimental design variables $\{z_i\}$, e.g., different loading paths, test specimen shapes, and observations. In this case, we have a set of data random vectors $\{Y_i(z_i)\}$ defined by

$$Y_i(z_i) = \mathcal{F}(\Theta, z_i, \eta_i), \quad \Theta \sim \pi_{\Theta}, \quad \eta_i \stackrel{\text{i.i.d.}}{\sim} \pi_{\eta_i}, \quad 1 \leq i \leq N_E.$$

In this case, the likelihood is given by

$$\pi_{\{Y_i(z_i)\}|\theta}(\{y_i\} | \theta) := \prod_{i=1}^{N_E} \pi_{Y_i(z_i)|\theta}(y_i | \theta). \tag{9}$$

We have the following Bayes' rule for the posterior of the constitutive parameters:

$$\pi_{\Theta|\{Y_i(z_i)\}}(\theta|\{y_i\}) = \frac{\pi_{\{Y_i(z_i)\}|\theta}(\{y_i\} | \theta) \pi_{\Theta}(\theta)}{\pi_{\{Y_i(z_i)\}}(\{y_i\})}$$

where the data marginal $\pi_{\{Y_i(z_i)\}}$ is given by

$$\pi_{\{Y_i(z_i)\}}(y) = \mathbb{E}_{\theta \sim \pi_{\Theta}} [\pi_{\{Y_i(z_i)\}|\Theta}(y | \theta)].$$

While we are interested in the BIP for multiple experiments, it is a straightforward extension of the BIP for a single experiment. We will use the simplified notation for a single experiment in the subsequent analysis, except when the distinction is important.

4. Bayesian optimal experimental design

Building on the BIP formulation from Section 3, Section 4.1 formulates Bayesian optimal experimental design (BOED) using the expected information gain (EIG), which covers (CI). Then, in Section 4.2, we present two equivalent forms of the EIG based on expected entropy reduction and mutual information. These formulations motivate the three EIG estimation approaches considered in this work:

- (i) Gaussian posterior approximation (Section 4.3), which is efficient but biased. This covers (CII).
- (ii) Surrogate Fisher information matrix (Section 4.4), which improves the efficiency of the Gaussian approximation for batched design. This covers (CIII).
- (iii) Nested Monte Carlo (Section 4.5), a baseline method that is asymptotically consistent but computationally demanding.

Lastly, in Section 4.6 we discuss the suitability of the first two approaches for learning history-dependent constitutive models from high-dimensional data.

4.1. Expected information gain as the design utility

We derive the EIG from a BIP perspective. Suppose that the data y is observed from the experiment with design z , and the BIP is solved to obtain the posterior $\pi_{\Theta|Y(z)}(\cdot|y)$. The success of the experiment can be assessed via the reduction of uncertainty, or the *information gain* in the constitutive parameters. A useful metric for the level of uncertainty in a probability distribution p is given by the information entropy \mathcal{H} defined as

$$\mathcal{H}(p) := -\mathbb{E}_{x \sim p} [\log p(x)].$$

Following our Bayesian formulation of parameter identification in Section 3, the information gain (IG) from an experiment with the design z and observed data y can be defined through the difference in the information entropy between the prior and the posterior:

$$\text{IG}(y, z) := \mathcal{H}(\pi_{\Theta}) - \mathcal{H}(\pi_{\Theta|Y(z)}(\cdot|y)). \quad (10)$$

While the IG enables us to assess data quality for identifying constitutive parameters, the primary goal of BOED is to propose experiments without access to the experimental outcomes. A natural solution to this dilemma is to use the forward model F to simulate data. This gives rise to the expected IG (EIG) defined as an expectation of the IG over the data marginal $\pi_{Y(z)}$:

$$\text{EIG}(z) := \mathbb{E}_{y \sim \pi_{Y(z)}} [\text{IG}(y, z)]. \quad (11)$$

To see where the prior and the forward model come into play in EIG, notice that sampling from the data marginal $y \sim \pi_{Y(z)}$ follows a procedure given by:

- (i) Sample constitutive parameters from the prior $\theta \sim \pi_{\Theta}$.
- (ii) Sample noise $\eta \sim \pi_{\eta}$.
- (iii) Evaluate the forward model $F(\theta, z, \eta)$ to obtain a simulated data sample at the constitutive parameters θ and the experimental design z .

The objective of BOED is to find the optimal design z^{\dagger} that maximizes the EIG subject to regularization:

$$z^{\dagger} \in \underset{z \in \mathbb{R}^{d_z}}{\text{argmax}} \{ \text{EIG}(z) - \mathcal{P}(z) \}, \quad (12)$$

where the regularization function $\mathcal{P} : \mathbb{R}^{d_z} \rightarrow \overline{\mathbb{R}}$ imposes bound constraints and penalties on the design variables. In this section, we introduce three different estimators for the EIG. Once an estimator is chosen, the maximization in Eq. (12) can be solved using an appropriate optimization algorithm. Examples include Bayesian optimization, which is applicable when the estimator's gradient is unavailable, and stochastic gradient descent, which is applicable when an unbiased gradient estimator is available; see [5, Section 4] for a detailed discussion of these algorithms.

4.2. Equivalent forms of the expected information gain

The EIG has numerous equivalent forms and interpretations. While they have been discussed in the literature (see, e.g., [5,14,16]), we provide the following lemma to make an explicit connection between the expected entropy reduction form in Eq. (11) and the mutual information forms of the EIG. The proof can be found in Appendix B.

Lemma 4.1 (The Expected Information Gain as Mutual Information). *The EIG, defined as the expected reduction in information entropy from the prior to the posterior in Eq. (11), is equivalent to the mutual information, denoted by MI, between the parameters and the data:*

$$\text{EIG}(z) \equiv \text{MI}(\Theta; Y(z)).$$

Here, the mutual information has two equivalent definitions:

$$\text{MI}(\Theta; Y(z)) := \mathbb{E}_{y \sim \pi_{Y(z)}} \left[D_{\text{KL}} \left(\pi_{\Theta|Y(z)}(\cdot | y) \parallel \pi_{\Theta} \right) \right],$$

$$\text{MI}(\Theta; Y(z)) := \mathbb{E}_{\theta \sim \pi_{\Theta}} \left[D_{\text{KL}} \left(\pi_{Y(z)|\Theta}(\cdot | \theta) \parallel \pi_{Y(z)} \right) \right],$$

where D_{KL} denotes the Kullback–Leibler divergence.

While theoretically identical, the two forms of the EIG pose distinct computational challenges. The expected entropy reduction form is conceptually intuitive and directly measures uncertainty minimization, but it is computationally expensive for nonlinear models because it requires repeated posterior density evaluations and sampling across data realizations. The likelihood-based mutual information form is useful when likelihood evaluation is tractable, but estimating the marginal data density $\pi_{Y(z)}$ can be difficult. The remainder of this section is devoted to deriving estimators of the EIG starting from its two forms. Their advantages and limitations when learning history-dependent constitutive laws will be discussed.

4.3. Estimator based on Gaussian approximation

We consider the EIG in the form of Eq. (11). Estimating EIG in this form is often intractable due to the prohibitive cost of posterior sampling and density evaluations across many data realizations, each of which requires running MCMC and evaluating the data marginal density. This issue can be addressed by first constructing a Gaussian posterior approximation and then using it in place of the posterior for approximate, rapid EIG estimation.

This subsection is organized as follows. We derive a Gaussian approximation to the posterior in Section 4.3.1. The formulas for the resulting EIG approximation that yield the Bayesian D-optimal design utility are shown in Section 4.3.2. Lastly, we discuss the numerical evaluations as well as limitations and advantages of this approximation in Sections 4.3.3 and 4.3.4, respectively.

4.3.1. Gaussian posterior approximation

First, notice that the data realization y that is used to evaluate the IG in Eq. (11) is generated by forward model evaluation $y = \mathcal{F}(\theta, z, \eta)$ at the prior sample $\theta \sim \pi_{\Theta}$ and the noise sample $\eta \sim \pi_{\eta}$, as described in Section 4.1. We consider a data-independent local Gaussian approximation, denoted as $\pi_{\text{loc}}(\cdot; \theta, z)$. It captures the expected posterior behavior given the data-generating parameters θ , averaging out variation across data realizations due to the random noise η . Compared with the Laplace approximation, which centers the Gaussian approximation at the MAP point that depends on the data realization, this data-independent approximation circumvents the optimization step for parameter identification. This optimization procedure can be computationally intensive when evaluating the forward model requires solving a transient PDE with history-dependent constitutive models.

We make the following assumptions on the likelihood and the prior:

- (i) The log-likelihood $\ell(\theta'; y, z) := \log \pi_{Y(z)|\Theta}(y | \theta')$ is twice continuously differentiable.
- (ii) The likelihood satisfies regularity conditions so that differentiation may be interchanged with expectation under $y \sim \pi_{Y(z)|\Theta}(\cdot | \theta)$.
- (iii) The log-prior $\log \pi_{\Theta}$ is twice continuously differentiable on the interior of its support.
- (iv) The prior π_{Θ} is strictly log-concave on the interior of its support.

At the data-generating parameter θ , the expected score satisfies the first Bartlett identity [84], and the Fisher information matrix (FIM) is defined as the negative expected Hessian of the log-likelihood:

$$\mathbb{E}_{y \sim \pi_{Y(z)|\Theta}(\cdot|\theta)} \left[\nabla_{\theta'} \ell(\theta; y, z) \right] = 0, \quad \text{FIM}(\theta, z) := -\mathbb{E}_{y \sim \pi_{Y(z)|\Theta}(\cdot|\theta)} \left[\nabla_{\theta'}^2 \ell(\theta; y, z) \right]. \quad (13)$$

We define the expected log-posterior (up to an additive constant independent of θ') as

$$\phi(\theta'; \theta, z) := \mathbb{E}_{y \sim \pi_{Y(z)|\Theta}(\cdot|\theta)} \left[\ell(\theta'; y, z) \right] + \log \pi_{\Theta}(\theta').$$

Using Eq. (13), the gradient and Hessian of ϕ at θ are

$$\nabla_{\theta'} \phi(\theta; \theta, z) = \nabla \log \pi_{\Theta}(\theta), \quad \nabla_{\theta'}^2 \phi(\theta; \theta, z) = -\text{FIM}(\theta, z) + \nabla^2 \log \pi_{\Theta}(\theta).$$

Consequently, a quadratic expansion of ϕ around θ is given by

$$\begin{aligned} \phi_{\text{quad}}(\theta'; \theta, z) := & \phi(\theta; \theta, z) + (\theta' - \theta)^{\top} \nabla \log \pi_{\Theta}(\theta) \\ & - \frac{1}{2} (\theta' - \theta)^{\top} \left(\text{FIM}(\theta, z) - \nabla^2 \log \pi_{\Theta}(\theta) \right) (\theta' - \theta), \end{aligned}$$

where the matrix $\text{FIM}(\theta, z) - \nabla^2 \log \pi_{\Theta}(\theta)$ is the negative Hessian of ϕ at θ . We define the corresponding local Gaussian approximation

$$\pi_{\text{loc}}(\theta'; \theta, z) \propto \exp(\phi_{\text{quad}}(\theta'; \theta, z)).$$

Completing the square for ϕ_{quad} gives us the local Gaussian approximation

$$\pi_{\text{loc}}(\cdot; \theta, z) = \mathcal{N}(m_{\text{loc}}(\theta, z), \Gamma_{\text{loc}}(\theta, z)), \tag{14a}$$

$$m_{\text{loc}}(\theta, z) = \theta + \Gamma_{\text{loc}}(\theta, z) \nabla \log \pi_{\theta}(\theta), \quad \Gamma_{\text{loc}}(\theta, z) = \left(\text{FIM}(\theta, z) - \nabla^2 \log \pi_{\theta}(\theta) \right)^{-1}. \tag{14b}$$

4.3.2. Applying the Gaussian approximation

The approximate IG using Eq. (14), denoted as $\widetilde{\text{IG}}$, is given by

$$\text{IG}(y, z) \approx \widetilde{\text{IG}}(\theta, z) := \mathcal{H}(\pi_{\theta}) - \mathcal{H}(\pi_{\text{loc}}(\cdot; \theta, z)), \quad y \sim \pi_{Y(z)|\theta}(\cdot | \theta), \tag{15}$$

where θ is the generating parameter of the data y . Since the information entropy of Gaussian distributions has a closed-form expression, we have the following formula for the approximate IG in Eq. (15):

$$\widetilde{\text{IG}}(\theta, z) = \frac{1}{2} \log \det (\text{FIM}(\theta, z) - \nabla^2 \log \pi_{\theta}(\theta)) + \mathcal{H}(\pi_{\theta}), \tag{16a}$$

where the first term measures the volume of the posterior confidence ellipsoid.

Using the approximate IG, we obtain an approximation to the EIG:

$$\text{EIG}(z) = \mathbb{E}_{\theta \sim \pi_{\theta}} \left[\mathbb{E}_{y \sim \pi_{Y(z)|\theta}(\cdot | \theta)} [\text{IG}(y, z)] \right] \approx \mathbb{E}_{\theta \sim \pi_{\theta}} [\widetilde{\text{IG}}(\theta, z)], \tag{17}$$

where the expectation with respect to the data y collapses after the approximation. The EIG approximation in Eq. (17) accounts for global parametric uncertainty by prior expectation and is often referred to as the utility function for *Bayesian D-optimal design* [4]. This utility can be estimated using the sample average approximation, which leads to the *estimator based on Gaussian approximation* (EGA):

$$\begin{aligned} \theta^{(i)} &\stackrel{\text{i.i.d.}}{\sim} \pi_{\theta}, \quad 1 \leq i \leq N_P, \\ \text{EGA}(z; N_P) &:= \frac{1}{2N_P} \sum_{i=1}^{N_P} \log \det (\text{FIM}(\theta^{(i)}, z) - \nabla^2 \log \pi_{\theta}(\theta^{(i)})) + \mathcal{H}(\pi_{\theta}). \end{aligned} \tag{18}$$

Remark 1 (Gaussian Prior as a Special Case). When the prior is Gaussian with mean $m_0 \in \mathbb{R}^{d_{\theta}}$ and covariance $\Gamma_0 \in \mathbb{R}^{d_{\theta} \times d_{\theta}}$, we have

$$m_{\text{loc}} = \theta - \Gamma_{\text{loc}} \Gamma_0^{-1} (\theta - m_0), \quad \Gamma_{\text{loc}} = (\text{FIM}(\theta, z) + \Gamma_0^{-1})^{-1}.$$

Consequently, the EGA becomes:

$$\text{EGA}(z; N_P) := \frac{1}{2N_P} \sum_{i=1}^{N_P} \log \det (\text{FIM}(\theta^{(i)}, z) + \Gamma_0^{-1}) + \frac{1}{2} \log \det (\Gamma_0).$$

This formula is used in our numerical case studies.

Remark 2 (Gaussian Approximation Breaks Form Equivalence). Under a Gaussian posterior approximation, the two equivalent definitions in Lemma 4.1 diverge. When the prior is also Gaussian, the KL divergence from the posterior to the prior has a closed-form expression that includes a log-determinant term, along with additional terms accounting for the averaged posterior variance and the mean shift. While this approximation is widely used [27,29,34], it differs from standard Bayesian D-optimal design and does not generally apply to non-Gaussian priors. Furthermore, this approach underestimates the EIG as evident by the Barber–Agakov bound [85]. Indeed, this specific property makes it useful for variational approximation of the EIG; see, e.g., [86–89].

4.3.3. Numerical evaluation

Here, we discuss the numerical evaluation of the EGA. First, the FIM is equivalent to the Gauss–Newton Hessian for likelihood based on certain exponential families of distributions. When the equivalence holds, the FIM can be evaluated directly from the forward or adjoint sensitivity of the forward model. For example, when the noise model is additive and Gaussian, i.e., $y = \mathcal{F}_{\text{pto}}(\theta, z) + \eta$ with $\eta \sim \mathcal{N}(0, \Gamma_{\eta})$, we have

$$\text{FIM}(\theta, z) = \partial_{\theta} \mathcal{F}_{\text{pto}}(\theta, z)^{\top} \Gamma_{\eta}^{-1} \partial_{\theta} \mathcal{F}_{\text{pto}}(\theta, z),$$

where $\partial_{\theta} \mathcal{F}_{\text{pto}}(\theta, z) \in \mathbb{R}^{d_y \times d_{\theta}}$ is the partial derivative of the observables, i.e., noise-free simulated data, with respect to the parameters. In this case, evaluating the FIM requires $\min(d_y, d_{\theta})$ evaluations of the direct or adjoint sensitivity of the forward problem, where each sensitivity computation represents a matrix–vector multiplication or transpose multiplication by $\partial_{\theta} \mathcal{F}_{\text{pto}}(\theta, z)$. Second, the log-determinant of the posterior covariance matrices in Eq. (16) can be evaluated using the eigenvalues of an eigendecomposition of the FIM. This means that $\widetilde{\text{IG}}$ can be evaluated without forming FIM using matrix-free eigendecomposition algorithms. One can exploit the fast eigenvalue decay of the FIM to remove the dependence of the number of forward and/or adjoint sensitivity computations on the parameter dimension [90]. It has been applied for D- and A-optimal design for parameters with arbitrarily high dimensions [15,29,91,92].

4.3.4. Limitations and advantages

The Bayesian D-optimal design utility approximates the EIG effectively when posteriors are unimodal and locally Gaussian, as supported by the Bernstein–von Mises theorem [93, Chapter 4] in the large-data or small noise limit. Leveraging the asymptotic consistency of the MLE, we construct our local approximation via a quadratic expansion of the log-posterior around the data-generating parameter θ . However, this approximation requires that the FIM be non-singular and the likelihood be unimodal. It breaks down due to unidentifiability or multimodality arising from under-constrained models (e.g., interchangeable viscous branches) or from designs that produce uninformative data.

Despite the EGA's limited use cases, it often strikes a superior balance between design improvements and cost, as discussed in Section 4.6. Additionally, it enables interpretation of the experimental design via EGAs for marginal parameters, as extracting the marginal posteriors for the Gaussian approximation is straightforward. In our numerical examples, we use the EGA to interpret how experimental designs affect the identifiability of subsets of the constitutive parameters associated with different material behaviors.

4.4. Estimator based on surrogate Fisher information matrix

The computational cost of EGA-based design optimization can be overwhelming when designing multiple experiments at once (see Section 3.4), where EGA becomes

$$\text{EGA}(\{z_j\}; N_P) := \frac{1}{2N_P} \sum_{i=1}^{N_P} \log \det \left(\sum_{j=1}^{N_E} \text{FIM}(\theta^{(i)}, z_j) - \nabla^2 \log \pi_{\theta}(\theta^{(i)}) \right) + \mathcal{H}(\pi_{\theta}). \quad (19)$$

$\theta^{(i)} \stackrel{\text{i.i.d.}}{\sim} \pi_{\theta}, \quad 1 \leq i \leq N_P,$

Each evaluation of the EGA requires $N_P \times N_E$ evaluations of the forward model and $N_P \times N_E \times \min(d_{\theta}, d_y)$ linearized PDE solutions corresponding to the direct or adjoint sensitivity of the forward model. Repeated evaluations of FIM at different constitutive parameters and design variables become a bottleneck in the design optimization. To remove this bottleneck, we use a fast-to-evaluate neural network surrogate of the FIM, denoted as $\widetilde{\text{FIM}}(\theta, z; w)$ with trainable weights $w \in \mathbb{R}^{d_w}$, to replace the FIM evaluations in EGA. We refer to the resulting surrogate approximation of the EGA as the *estimator based on surrogate FIM* (ESFIM).

The surrogate is formulated to accommodate key features of FIMs, namely symmetric positive semi-definiteness and potentially extreme spectrum spreads. We consider using a neural network to emulate the lower-triangular components of the logarithmic regularized FIM:

$$B(\theta, z) := \log \text{m}(\text{FIM}(\theta, z) + \epsilon I), \quad \epsilon > 0,$$

$$\text{NN}(\theta, z; w) \approx \begin{bmatrix} B_{11} & B_{22} & \dots & B_{d_{\theta}d_{\theta}} & \sqrt{2}B_{12} & \sqrt{2}B_{13} & \dots \end{bmatrix}^{\top},$$

where $\log \text{m}(\cdot)$ is the matrix logarithm, ϵ is a small regularization constant, and $\text{NN} : \mathbb{R}^{d_{\theta}} \times \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_{\theta}(d_{\theta}+1)/2}$ is the neural network predictions of the components of $B(\theta, z)$. The scaling by $\sqrt{2}$ of the off-diagonal term is introduced to be consistent with the scaling in the Frobenius matrix norm. The neural network can be trained to improve its FIM prediction quality, for example, by using the mean relative Frobenius norm of the prediction error as the loss function. The training data can be generated by evaluating the FIM at samples of constitutive parameters and design variables. Once the neural network is trained, it can be used to predict FIMs:

$$\begin{bmatrix} \widetilde{B}_{11} & \widetilde{B}_{22} & \dots & \widetilde{B}_{d_{\theta}d_{\theta}} & \sqrt{2}\widetilde{B}_{12} & \sqrt{2}\widetilde{B}_{13} & \dots \end{bmatrix} := \text{NN}(\theta, z; w)^{\top},$$

$$\widetilde{\text{FIM}}(\theta, z; w) := \text{expm}(\widetilde{B}(\theta, z; w)),$$

where $\text{expm}(\cdot)$ is the matrix exponential. The surrogate predictions of the FIM are then used in the evaluations of ESFIM by replacing FIM with $\widetilde{\text{FIM}}$ in Eq. (19), and subsequently in Eqs. (16) and (18).

Since the surrogate FIM is trained on random designs, the optimized designs found by ESFIM are likely to outperform random designs and perform similarly or worse than those found by EGA. The dominant computational cost of ESFIM-based optimal design arises from surrogate construction, which is independent of the number of experiments, N_E . Consequently, ESFIM enables optimization of experimental design for a large number of experiments, achieving a greater reduction in parametric uncertainty than random designs at a much lower computational cost than EGA. This advantage of ESFIM is demonstrated in our numerical examples.

4.5. Nested Monte Carlo estimator

Lastly, we introduce the baseline method for EIG estimation: the nested Monte Carlo estimator. When the likelihood $\pi_{Y(z)|\theta}(\cdot | \theta)$ is tractable to evaluate and generative, i.e., sampling $y \sim \pi_{Y(z)|\theta}(\cdot | \theta)$ is straightforward, one often resorts to the likelihood-based mutual information form of the EIG as in Lemma 4.1:

$$\text{EIG}(z) = \mathbb{E}_{\theta \sim \pi_{\theta}} \left[\mathbb{E}_{y \sim \pi_{Y(z)|\theta}(\cdot | \theta)} \left[\log \left(\frac{\pi_{Y(z)|\theta}(y | \theta)}{\pi_{Y(z)}(y)} \right) \right] \right]. \quad (20)$$

Here, the data marginal distribution $\pi_{Y(z)}$ is also defined through an expectation as in Eq. (8). Using the Monte Carlo (MC) method to estimate the expectations in Eq. (20), we arrive at a nested Monte Carlo estimator (NMCE) [32] given by:

$$\theta^{(i)} \stackrel{\text{i.i.d.}}{\sim} \pi_{\theta}, \quad y^{(i)} \stackrel{\text{i.i.d.}}{\sim} \pi_{Y(z)|\theta}(\cdot | \theta^{(i)}), \quad \theta^{(ij)} \stackrel{\text{i.i.d.}}{\sim} \pi_{\theta}, \quad 1 \leq i \leq N_O, \quad 1 \leq j \leq N_I, \tag{21}$$

$$\text{NMCE}(z; N_O, N_I) := \frac{1}{N_O} \sum_{i=1}^{N_O} \log \left(\frac{\pi_{Y(z)|\theta}(y^{(i)} | \theta^{(i)})}{\frac{1}{N_I} \sum_{j=1}^{N_I} \pi_{Y(z)|\theta}(y^{(i)} | \theta^{(ij)})} \right).$$

The NMCE is biased for a finite N_I but is asymptotically unbiased and consistent [33, Theorem 1, Appendix B]. The computational effort required to achieve a mean squared error of ϵ^2 is shown to be $O(\epsilon^{-3})$ for continuously differentiable log-likelihood, which is slower than the standard MC rate of $O(\epsilon^{-2})$ [33, Appendix G]. Each evaluation of the NMCE requires $N_O(N_I + 1)$ solutions of the balance equation in Section 3.1, and achieving a desirable level of error tolerance in the EIG estimation can be prohibitively costly. Typically, the number of inner MC samples, N_I , needs to be large because the likelihood can be concentrated when the data are informative. Additional techniques, such as importance sampling, are often necessary when evaluating the NMCE; see [5, Section 3.1].

4.6. Discussion

We close this section by highlighting two central aspects of BOED, for learning history-dependent constitutive models, that make the EGA and ESFIM more robust and practical than the NMCE baseline in design optimization.

Cost-accuracy trade off. For history-dependent constitutive models, the dominant cost is the repeated solution of the balance equation over the loading history. In this setting, NMCE can be prohibitively expensive because it places this cost inside a nested Monte Carlo loop. EGA removes this bottleneck by replacing this nested loop with a single loop of FIM evaluations, and ESFIM bypasses this cost entirely after a surrogate FIM is trained. It remains to be shown that EGA and ESFIM produce good designs despite the bias they introduce into the EIG estimation, and this is addressed empirically through numerical examples in the next two sections.

High-dimensional data. High-resolution images can extract large amounts of information relevant to parameter identification, provided that the design excites the relevant heterogeneous responses. EGA and ESFIM remain viable in this regime because their dominant costs, namely linearized PDE solves and surrogate training, are independent of the data dimension. By contrast, Section 5.4 of our numerical example shows that the NMCE baseline is intractable to evaluate even with importance sampling due to extreme likelihood concentration induced by the image data.

5. Application: Learning linear viscoelasticity

In this section, we present a numerical study on designing uniaxial tests, as described in Section 3.1, to learn the constitutive parameters of the linear viscoelastic model described in Section 2.2. First, in Section 5.1 we specify the design problem’s settings, including the prior and observation operators. We then present results on design optimization for a single experiment and for a batch of experiments in Sections 5.2 and 5.3, respectively, which are overviewed as follows.

- (i) Surrogate FIM sample complexity that reaches below 3% relative error in log-FIM.
- (ii) Visualizations of the EGA- and ESFIM-optimized designs with load/unload patterns and tilted elliptical holes that agree with physical intuition about designs that maximize information.
- (iii) Quantitative interpretations of the optimized designs compared to random designs through credible interval size (averaged 47% uncertainty reduction) and marginal EGA values (10% improvement in learning relaxation times).
- (iv) A Spearman correlation of 0.67 between EGA and stress-state entropy, indicating that designs favored by EGA tend to induce more diverse stress states, but not strongly enough for the latter to serve as a surrogate utility.

Finally, in Section 5.4 we demonstrate the intractability of NMCE and the suitability of EGA for this design problem. The implementation details are provided in Appendix C.

5.1. Settings

Experimental design. The uniaxial testing setup described in Section 3.1 is used with the following specifications shown in Fig. 3. For shape design, we consider a specimen with normalized widths $L_1 = 2$ and $L_2 = 1$. The hole Ω_H in the specimen has an elliptical shape, parameterized by one of the principal angles $\alpha_s \in [0, \pi/2]$. The principal axis lengths, denoted as $L_{s,1}$ and $L_{s,2}$, are each chosen in the interval $[0.1, 0.35]$. Together, there are three shape design variables $z_s = (\alpha_s, L_{s,1}, L_{s,2})$. For the loading path design, a normalized total test time of $T = 1$ is considered, and 10 loading path control values are used. Specifically, the 10 components of $z_l \in [0, 0.1]^{10}$ are taken as loading values on an equally spaced grid of 10 points in time-interval $[0.1, 1.0]$, and interpolated, starting from zero loading at time $t = 0$. The interpolation is performed using a monotonic piecewise cubic Hermite interpolating polynomial (PCHIP). The choice of range of z_l leads to a maximum imposed strain of 5%. In Fig. 3, this experimental setup and the design variables are visualized. \diamond

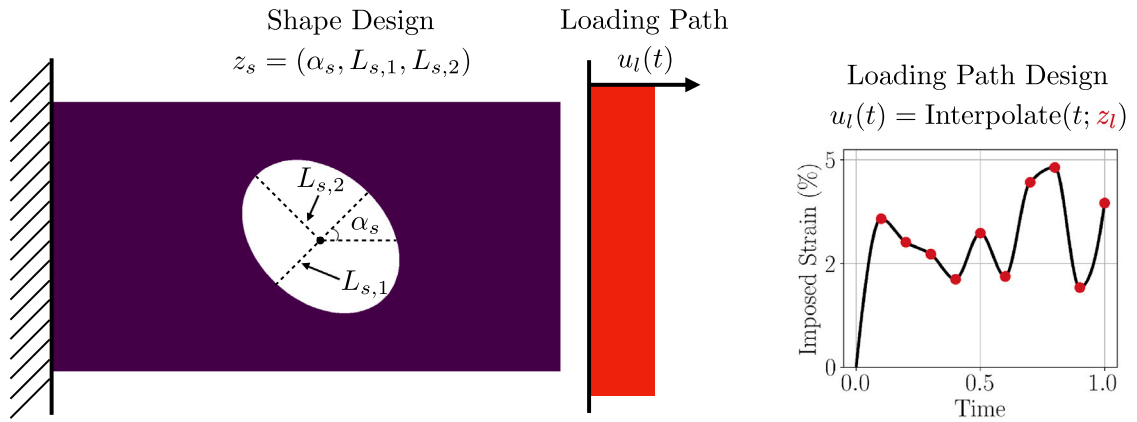


Fig. 3. An example of the uniaxial testing setup and the design variables. (Left) The setup and the specimen geometry design. (Right) The loading path design.

Table 1

The bounds for the uniform distributions of the physical parameters in the re-parameterized linear viscoelastic constitutive model.

Physical parameter	Description	Bound
$\log E_1^{(0)}$	Log. of equilibrium modulus in 1-dir.	(2.0, 3.5)
r_E	Equilibrium modulus anisotropy, E_2/E_1	(0.2, 1)
r_G	Normalized shear, $G/\sqrt{E_1 E_2}$	(0.2, 0.8)
r_ν	Poisson's ratio product, $\nu_{12}\nu_{21}$	(0.0, 0.2)
α_c	Material axis angle	$(-\pi/4, \pi/4)$
$f_i, i = 1, 2$	Total relaxed fraction in dir. i	(0.2, 0.8)
$w_j^{(1)}, j = 1, 2$	Branch-1 share of f_i	(0, 1)
$\log \tau_1$	Log. of relaxation time, fast	(-3.4, -1.2)
$\log \tau_2$	Log. of relaxation time, slow	(-1.2, 1)

Reparameterization and the prior distribution. We consider the linear viscoelastic constitutive model in Section 2.2, simplified under the orthotropic and plane-stress assumptions with two viscous branches. This reparameterization consists of two steps. First, a standard normal prior distribution, i.e., $\pi_\theta = \mathcal{N}(0, I)$, is prescribed. Then, the physical parameters, denoted as θ , are obtained via the following transformation

$$\theta_i = a_i + \frac{(b_i - a_i)}{2} \left(1 + \operatorname{erf} \left(\frac{\theta_i}{\sqrt{2}} \right) \right), \quad \theta_i \sim \mathcal{U}(a_i, b_i), \quad (22)$$

where $\{a_i\}$ and $\{b_i\}$ are the lower and upper bounds of the uniform distributions of the physical parameters. The physical parameters and their ranges are specified in Table 1.

Here, we consider a parametrization for the elastic and viscous tensors $\{\mathbb{C}_i\}$ and justify the associated ranges. Throughout, Voigt notation uses the in-plane ordering [11, 22, 12], and a single material orientation is shared across all branches:

$$\mathbb{C}_i = \mathbb{T}(\alpha_c) \mathbb{S}_i^{-1} \mathbb{T}(\alpha_c)^\top, \quad \mathbb{S}_i = \begin{bmatrix} 1/E_1^{(i)} & \nu_{21}^{(i)}/E_2^{(i)} & 0 \\ \nu_{12}^{(i)}/E_1^{(i)} & 1/E_2^{(i)} & 0 \\ 0 & 0 & 1/G_{12}^{(i)} \end{bmatrix}, \quad 0 \leq i \leq 2,$$

where $\mathbb{T}(\alpha_c)$ is the in-plane rotation by angle α_c , and \mathbb{S}_i denotes the material-frame plane-stress orthotropic compliance with four independent components per branch among $E_1^{(i)}$, $E_2^{(i)}$, $G_{12}^{(i)}$, $\nu_{12}^{(i)}$ and $\nu_{21}^{(i)}$. The material axis orientation is uniformly distributed as $\alpha_c \sim \mathcal{U}(-\pi/4, \pi/4)$. This range of orientations is chosen because we hypothesize that the optimal ellipse is likely tilted, which creates an interesting experimental design problem.

We set uniform priors for the logarithm of the modulus $E_1^{(0)}$ in the equilibrium branch and the ratio between $E_2^{(0)}$ and $E_1^{(0)}$,

$$\log E_1^{(0)} \sim \mathcal{U}(2, 3.5), \quad r_E \sim \mathcal{U}(0.2, 1), \quad E_2^{(0)} := r_E E_1^{(0)}.$$

The range of $E_1^{(0)}$ is chosen so that the reaction force in Eq. (23) is, on average, of order 1. The viscous modulus $E_1^{(i)}$ and $E_2^{(i)}$ for $i = 1, 2$ are related to the equilibrium ones via the total fractions f_1, f_2 :

$$f_j \sim \mathcal{U}(0.2, 0.8), \quad E_j^{(i)} = f_j w_j^{(i)} E_j^{(0)}, \quad 1 \leq i, j \leq 2,$$




Fig. 4. An example of simulated force and image data for learning linear viscoelasticity. This simulated dataset is created via the experimental setup in Section 5.1 and Fig. 3, and the constitutive parameters are drawn from the prior. Note that the resolution of the image data visualized here is much lower than that of the actual data used in the numerical study; see an image snapshot with high resolution in Fig. C.17. The material deformation is discernible yet small due to the narrow range of the imposed strain ($< 5\%$) suitable for linear viscoelasticity.

where the weights are given by

$$w_j^{(1)} \sim \mathcal{U}(0, 1), \quad w_j^{(2)} := 1 - w_j^{(1)}, \quad 1 \leq j \leq 2.$$

For all branches, the shear moduli $\{G^{(i)}\}$ are defined through a single ratio r_G

$$r_G := G_{12}^{(i)} / \sqrt{E_1^{(i)} E_2^{(i)}} \sim \mathcal{U}(0.2, 0.8).$$

The Poisson's ratios $\{\nu_{12}^{(i)}\}$ and $\{\nu_{21}^{(i)}\}$ are determined by a single ratio r_ν :

$$r_\nu := \nu_{12}^{(i)} \nu_{21}^{(i)} \sim \mathcal{U}(0, 0.2), \quad \nu_{12}^{(i)} := \sqrt{r_\nu E_1^{(i)} / E_2^{(i)}}, \quad \nu_{21}^{(i)} := \sqrt{r_\nu E_2^{(i)} / E_1^{(i)}}.$$

For this case study, the relaxation times are set to distinct scales, i.e., the time between control points (~ 0.1) and the experiment time window (~ 1), to ensure identifiability of the branches:

$$\log \tau_1 \sim \mathcal{U}(-3.4, -1.2), \quad \log \tau_2 \sim \mathcal{U}(-1.2, 1).$$

◇

Observation operator. We use a combination of image data and force data for our numerical study. For image data, 20 snapshots are taken, equally spaced in time. We assume a pixel density of 500, with a field of view of height 1.06 and width 2.1, resulting in an image size of 530×1050 (height \times width). A fixed speckle pattern is used across different experiments, with an averaged speckle radius of 6×10^{-3} (6 pixels across a speckle on average), a contrast of 0.8, and coverage of approximately 50%. We assume $\sigma_\eta = 0.02$. The image observation operator is defined in Section 3.2 and detailed in Appendix A, and the image generation procedure is detailed in Appendix C.2. For the force data, 100 values are collected, equally spaced in time. The force data at time t_i is modeled as the outward force at the right boundary where the loading is applied:

$$y_{f,i} = \int_{\{2\} \times [0,1]} [P(X, t_i)]_{11} dX + \eta_f, \quad \eta_f \sim \mathcal{N}(0, 2.5^{-5}), \quad (23)$$

where a signal-to-noise ratio of about 0.5% is used. An example of simulated experimental data is shown in Fig. 4. ◇

5.2. Designing a single experiment

We maximize EGA and ESFIM to obtain designs of a single experiment. Here, we briefly describe the procedures for EGA evaluation, surrogate FIM training, and utility maximization. Implementation details are provided in Appendix C

EGA evaluation. At each parameter sample and design candidate, we compute the Jacobian of the parameter-to-observable map, $\partial_\theta \mathcal{F}_{\text{p10}}$ using forward sensitivity solves. This requires 11 linear, time-dependent PDE solves per parameter sample at each design candidate, in addition to the one required for forward-model evaluations.

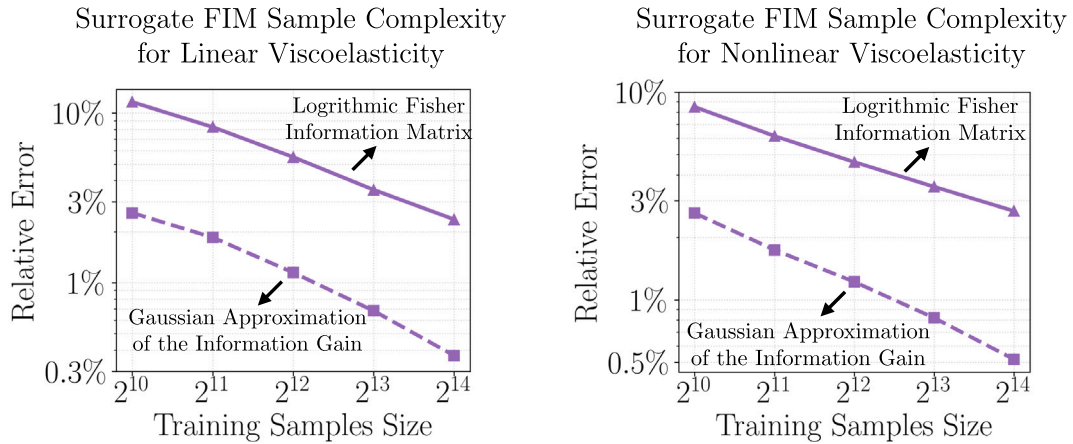


Fig. 5. The surrogate FIM testing accuracy as a function of training sample size for designing uniaxial tests of linear and nonlinear viscoelastic materials in Sections 5 and 6, respectively. The solid and dashed lines indicate the mean relative errors for predicting the log. FIM (measured in the Frobenius norm) and the approximate information gain, respectively. The surrogate prediction of the approximate information gain is obtained from the surrogate FIM predictions using Eq. (16).

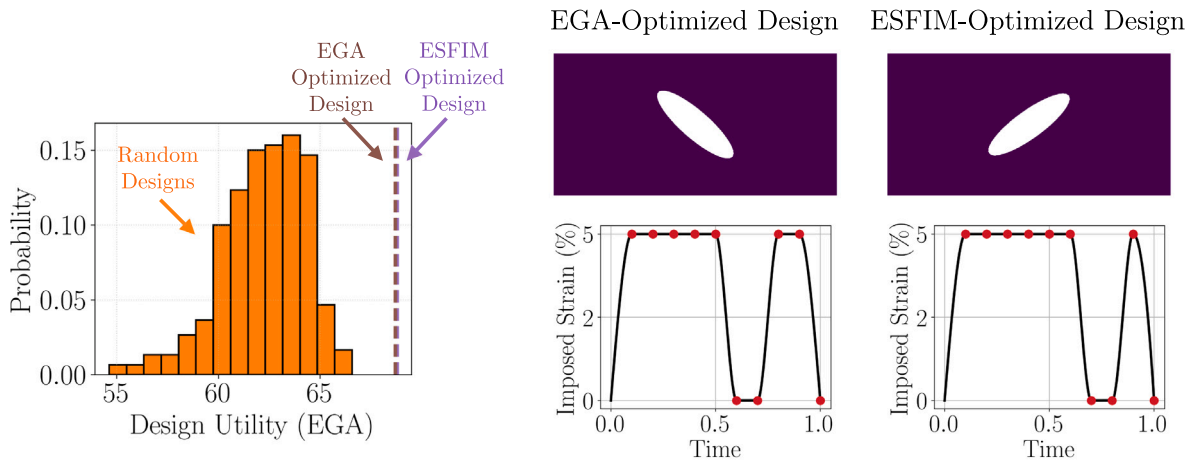
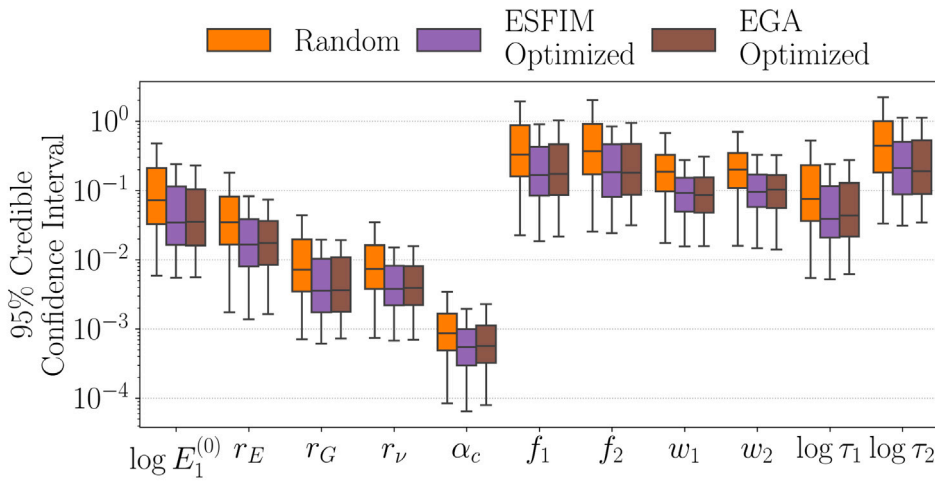


Fig. 6. The optimized designs of uniaxial test for reliable learning of linear viscoelasticity found by EGA and ESFIM maximization. (Left) Comparison of the EGA values at random, EGA-optimized, and ESFIM-optimized designs. (Right) Visualization of the optimized designs found via EGA maximization and ESFIM maximization.

Surrogate FIM. The surrogate model is a residual network with 6 hidden layers and Gaussian Error Linear Unit (GELU) activation, with a layer width of 200. The surrogate is trained on data consisting of FIMs computed at 16,384 parameter and design variable samples generated via QMC. The surrogate achieved a mean relative accuracy of over 97% in predicting the log. FIM on a testing dataset of 4096 samples. Moreover, the surrogate achieves a mean relative accuracy of over 99.6% in predicting the approximate information gain in Eq. (16) on the test set. On the left panel of Fig. 5, we visualize the surrogate FIM sample complexity, i.e., testing error as a function of training sample size. A standard MC rate of convergence at $N_{\text{train}}^{-1/2}$ is observed. \diamond

Design optimization. We use EGA and ESFIM to approximate the EIG in the design optimization of Eq. (12), subject to a bound constraint on the design variables and without additional regularization. The EGA is maximized using the Bayesian optimization method with $N_p = 128$. The ESFIM is maximized using L-BFGS-B with an ensemble of initial starts. We note that the data generation and training costs for constructing the surrogate FIM model are similar to the total costs of Bayesian optimization for EGA maximization. The ESFIM maximization costs on a Graphics Processing Unit (GPU) are relatively negligible compared to those of EGA maximization. \diamond

Now we present and analyze the EGA- and ESFIM-optimized designs, both qualitatively and quantitatively.



Expected Reduction of the 95% Credible Intervals By Optimized Designs Relative to Random Designs												
	$\log E_1^{(0)}$	r_E	r_G	r_ν	α_c	f_1	f_2	w_1	w_2	$\log \tau_1$	$\log \tau_2$	Averaged
ESFIM	55%	50%	48%	33%	48%	47%	44%	55%	51%	48%	42%	47%
EGA	54%	50%	50%	34%	52%	47%	44%	54%	51%	49%	44%	48%

Fig. 7. Visualizations and statistics of the 95% CIs for constitutive parameters of linear viscoelasticity inferred from experiments with random, ESFIM-optimized, and EGA-optimized designs. (Top) The box plots of the 95% CI sizes. The size distributions arise from the prior-based variation in the experimental outcomes. For random designs, the distributions also account for design variations. (Bottom) The relative expected reduction of the 95% CI sizes by optimized designs relative to random designs. The 95% CI sizes are estimated for θ , which is transformed as in Section 5.1 from the physical parameters used as labels.

5.2.1. Optimized designs

The EGA-optimized design is shown in Fig. 6. The loading path increases to the maximum strain at the first control point, holds for 6 consecutive control points, then rapidly unloads and holds for 2 control points, and then rapidly loads to the maximum strain and holds for 2 control points. The rapid loading and unloading, along with long holding periods, can be interpreted as maintaining a balance between learning the two relaxation times. The optimized shape of the elliptical hole has a maximum aspect ratio with $L_{s,1} = 0.1$ and $L_{s,2} = 0.35$, and an angle of $\alpha_s = 0.27\pi$ (48.9°). This tilted shape is expected due to the prior on the material axis angle specified in Table 1.

The ESFIM-optimized design is also shown in Fig. 6. The loading path increases to the maximum strain at the first control point, holds for 6 consecutive control points, then rapidly unloads and holds for 2 consecutive control points, and then rapidly loads and unloads for the final two control points. The optimized shape of the ellipse hole has $L_{s,1} = 0.35$ and $L_{s,2} = 0.1$, and an angle of $\alpha_s = 0.2\pi$ (35.9°). The ESFIM-optimized design can be interpreted similarly to the EGA-optimized design, with notable variations in hole orientation and holding duration. However, the resulting discrepancy between the two designs has a negligible effect on parameter identification, as indicated by their EGA values, which differ by less than 0.2%.

We compare optimized designs with an ensemble of 300 designs randomly sampled from the uniform distribution within the design variable’s bounds. The distribution of the EGA utility values for random designs is shown on the left panel of Fig. 6, along with the utility values of the EGA- and ESFIM-optimized designs. Compared to random designs, both the EGA- and ESFIM-optimized designs yield, on average, 10% higher EGA utility values. We note that the same set of random designs is used for the subsequent results in this section.

5.2.2. Quantitative interpretation: Credible intervals of the posteriors

In Fig. 7, we present the visualizations and statistics of the 95% credible intervals (CIs) for the posteriors. The credible intervals are computed using the Gaussian approximation in Section 4.3.1. For each experimental design, we use 128 samples of the data-generating parameters, drawn via QMC, to capture the variation in the experimental outcomes. In the top panel of Fig. 7, we visualize

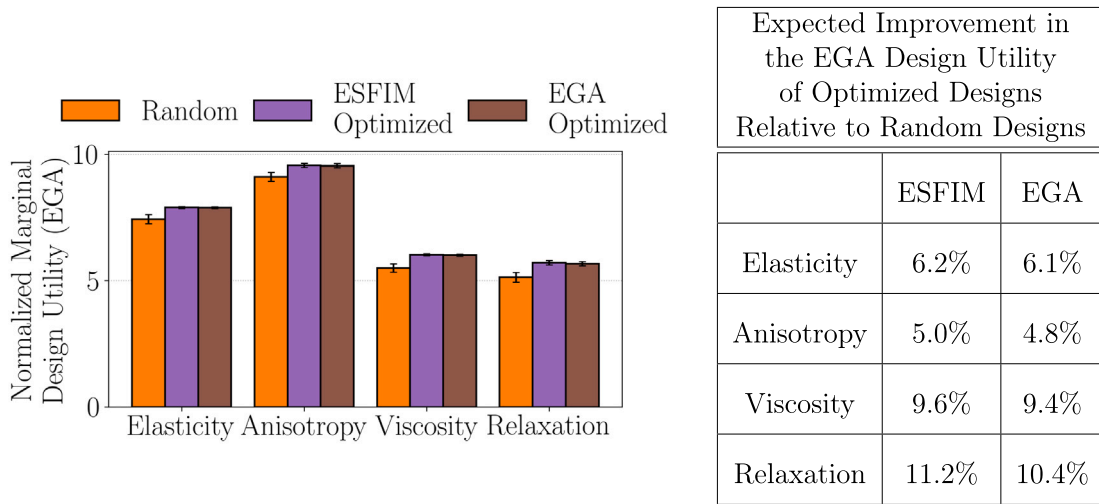


Fig. 8. The normalized EGA design utility for learning marginal parameters of linear viscoelasticity grouped by material behaviors. (Left) The bar plots of the utility of random, ESFIM-optimized, and EGA-optimized designs for learning elasticity, anisotropy, viscosity, and relaxation. The error bars for the random designs represent the standard deviations of EGA with respect to the design variables. In contrast, the error bars for the optimized designs represent the standard errors of EGA attributed to finite parameter sample sizes. (Right) The expected improvement in the utility of optimized design relative to random designs for learning different behaviors.

the distribution of 95% CI sizes across prior-based variation in experimental outcomes. For random designs, the size distribution also accounts for design variations. We observe that the optimized designs consistently shift the distribution of the 95% CI sizes toward lower values. In the bottom panel of Fig. 7, we provide the relative expected reduction of the 95% CIs by optimized designs compared to random designs. The formula is given as follows:

$$\frac{\sum_{i,j} (\text{CI}(\theta^{(i)}, z^{(j)}) - \text{CI}(\theta^{(i)}, z^{\dagger}))}{\sum_{i,j} \text{CI}(\theta^{(i)}, z^{(j)})} \times 100\%,$$

where CI corresponds to the estimated 95% CI size for the parameter of interests, $\{\theta^{(i)}\}$ are samples of the data-generating parameters, $\{z^{(j)}\}$ are random design samples, and z^{\dagger} is the optimized design. The table shows that the design optimization significantly reduces the 95% CI sizes, with an average reduction of 47%. These results demonstrate that our optimized designs yield more reliable parameter estimation across a range of materials whose constitutive parameters are modeled by the prior.

5.2.3. Quantitative interpretation: Design utility for learning different behaviors

In Fig. 8, we provide the EGA for inferring four disjoint subsets of parameters. The four subsets corresponding to different modes of material response: (i) elasticity, $(\log E_1^{(0)}, r_E, r_G, r_v)$, (ii) anisotropy, α_c , (iii) viscosity, (f_1, f_2, w_1, w_2) , and (iv) stress relaxation, $(\log \tau_1, \log \tau_2)$. The relative expected improvements in the EGA for marginal parameters are computed as follows:

$$\frac{\sum_j (\text{EGA}_{\text{marg}}(z^{(j)}) - \text{EGA}_{\text{marg}}(z^{\dagger}))}{\sum_j \text{EGA}_{\text{marg}}(z^{(j)})} \times 100\%,$$

where EGA_{marg} corresponds to the EGA for marginal parameters, $\{z^{(j)}\}$ are random design samples, and z^{\dagger} is the optimized design. The results show that the optimized designs consistently improve the identifiability of these parameter subsets compared to random designs. The relative improvements are most significant for the relaxation times at 11%, whereas anisotropy is the least significant, yet still at 5%.

Remark 3. Although the marginal EGA improvements in Fig. 8 are only 4%–12%, they are computed based on the log-determinant of the marginal posterior precision matrix, so even a modest improvement corresponds to a substantial reduction in marginal posterior uncertainty. For example, in Fig. 7, the optimized designs reduce the average 95% CI width of the orthotropic rotation parameter (α_c) by 52%, whereas in Fig. 8, the relative improvement, measured by the marginal EGA, is 4.8%. The improvements are not larger because the baseline EGA value is already high, which we attribute to the high-dimensional data with a low noise level.

5.2.4. Relation between information entropy and stress-state entropy

The utility of a mechanical testing design is often tied to the diversity of local strain histories and the stresses it induces, since parameter identifiability can be improved by richer mechanical responses. To examine this connection in our setting, we consider

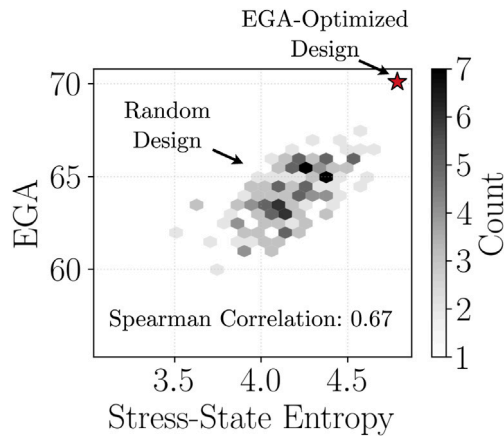


Fig. 9. Relation between principal-stress-based stress-state entropy and EGA for randomly selected designs and the EGA-optimized design. The EGA-optimized design also has the largest stress-state entropy, with an overall rank correlation of 0.67.

a principal-stress-based stress-state entropy [22,23]; its definition is given in Appendix C.3. Fig. 9 compares this mechanics-based quantity with EGA across random and optimized designs. The EGA-optimized design also attains the largest stress-state entropy, and the two quantities are positively correlated, with a Spearman rank correlation of 0.67. This indicates that designs favored by EGA tend to induce more diverse stress states, while also showing that stress-state entropy does not fully account for the information-based utility.

5.3. Batched experimental design

We utilize ESFIM with bound constraints on the design variables for batched optimal experimental design. Note that the same surrogate FIM described in Section 5.2 is reused to find all the optimizers shown below.

The optimizers for experiment batch sizes of 2 and 3 are shown in the top panel of Fig. 10. Within the same batch of experiments, the optimized shapes of the elliptical holes alternate in orientation, which is known to improve anisotropy identification. Loading paths with rapid loading and unloading, along with long holds, are still preferred by ESFIM, which is known to improve the identification of relaxation times. We observe an increase in the number of switches between no-strain and maximum-strain as the batch size increases, suggesting that improving the identifiability of the parameters associated with the instantaneous response becomes more rewarding.

In the bottom panel of Fig. 10, we visualize the optimized designs for batch sizes of 2 and 3 when specimen geometry is shared. This corresponds to reusing the same specimen across multiple loading paths, which is more practical for our small-strain setting. The result shows that the shared geometry for both batch sizes of 2 and 3 is a tilted ellipse with maximal aspect ratio: $L_{s,1} = 0.1$, $L_{s,2} = 0.35$, and $\alpha_s = 0.3\pi$ (54°). We again observe an increase in the number of switches between no strain and maximum strain as the batch size increases.

In Fig. 11, we visualize and quantify the improvements in the EGA for ESFIM-optimized designs shown in the top panel of Fig. 10 relative to random designs across different experiment batch sizes. For a single optimized experiment, the optimized design has a similar EGA utility value to that of three randomly designed experiments. We expect four randomly designed experiments to outperform one optimized experiment. On the other hand, at similar computational costs, ESFIM obtains optimized designs for experiment batch sizes of 1, 2, and 3, and each substantially outperforms the random designs of the same batch size by around 10%. This demonstrates ESFIM's cost-amortization capability and its suitability for batched design.

5.4. Nested Monte Carlo versus Gaussian approximation

We report and discuss our attempts to approximate the EIG using NMCE and EGA across three randomly sampled designs (see Fig. C.18). The outcomes suggest that NMCE, while asymptotically unbiased, is computationally intractable as a design utility due to the extreme likelihood concentration arising from high-resolution image data. The results motivate the use of alternative EIG approximation techniques, such as EGA, which strikes a superior balance between speed and accuracy compared to NMCE.

We use $N_{\text{out}} = 64$ and $N_{\text{in}} = 16,384$ for NMCE, with importance sampling using proposals given by the Gaussian approximation in Eq. (14). Both outer and inner sampling are performed with quasi-Monte Carlo (QMC) using scrambled Sobol sequences [94,95]. We observe severe weight collapse in the importance sampling. The effective sample size for inner sampling is 1 per outer sample. A variance decomposition of the log-weights shows that variation in the log-likelihood is the dominant contributor to the weight collapse. In particular, the log-likelihood of the data-generating parameter exceeds that of the inner samples by a median of 159, indicating an extreme concentration of likelihood that can be attributed to the data's high dimensionality. A log-normal

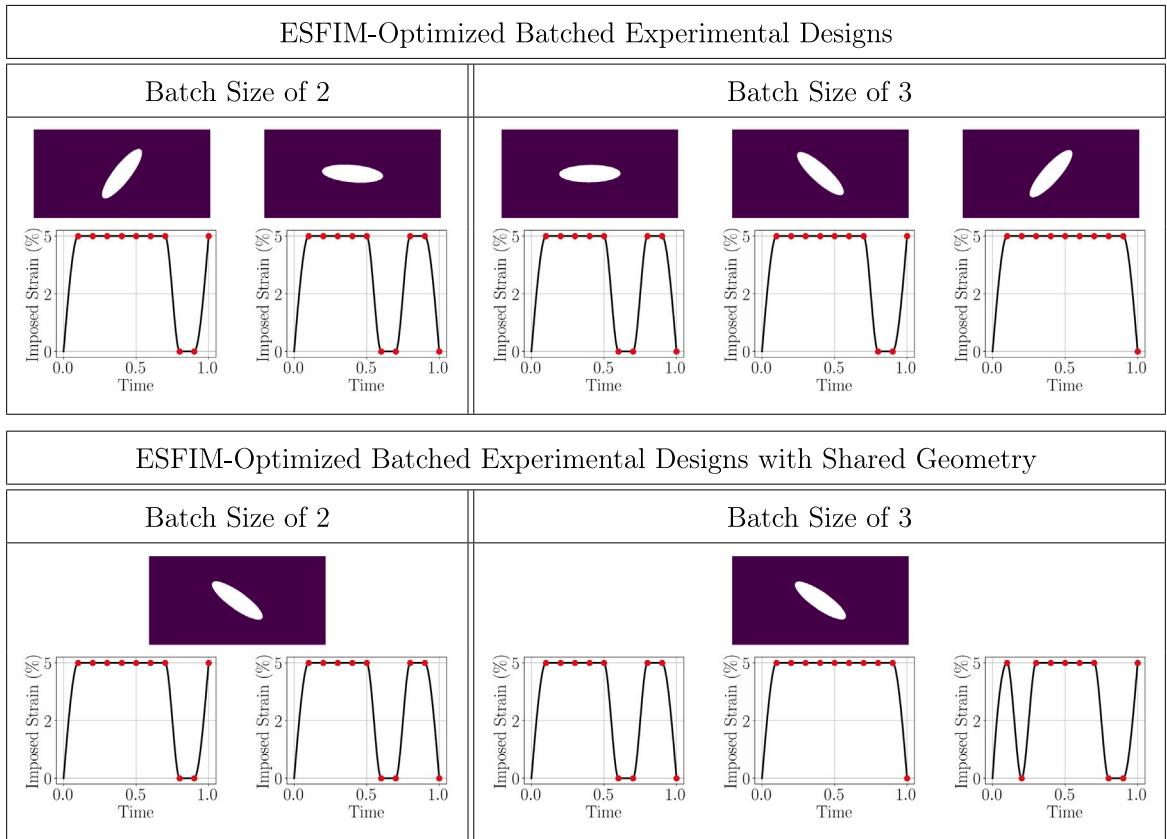
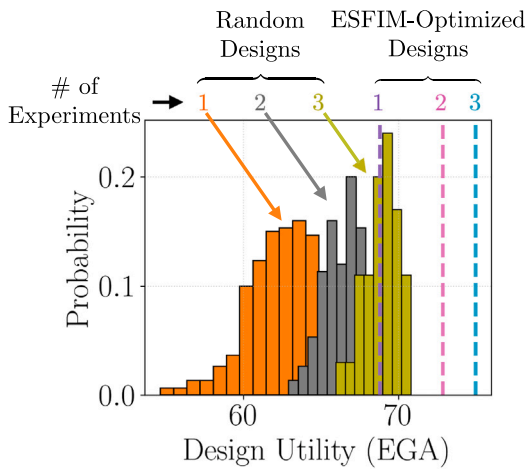


Fig. 10. The optimal batched designs of uniaxial testing for learning linear viscoelasticity found by ESFIM maximization. The designs in the top panel vary in geometry within their batches, whereas those in the bottom panel share the same geometry.



Expected Improvements in the EGA Design Utility of Optimized Designs Relative to Random Designs				
		Batch Size		
		1	2	3
Batch Size	Optimized	10.5%	3.4%	-0.2%
	1	17.0%	9.5%	5.7%
	2	20.4%	12.7%	8.8%

Fig. 11. Visualization and statistics of the EGA design utility values of batched designs for uniaxial testing of linear viscoelastic materials. (Left) The EGA values for random and ESFIM-optimized (batched) designs of 1–3 experiments. (Right) Relative expected improvements of the EGA for the optimized designs with batch sizes of 1–3 compared to random designs with batch sizes of 1–3.

approximation to the importance weights suggests that approximately 10^{10^7} samples would be required for accurate inner sampling, rendering NMCE intractable. See a more detailed discussion of this estimation in Appendix C.4.

In contrast, we observe that EGA evaluations, which share 64 outer samples with NMCE for each design, can be performed at a computational cost 4 orders of magnitude lower, while still yielding more reliable EIG approximations. Importantly, the costs of EGA evaluation are independent of the data dimension, whereas NMCE's efficiency and reliability degrade as the data dimension increases. Furthermore, this concentration of likelihood supports the validity of our Gaussian posterior approximation under the Bernstein-von Mises theorem; see Section 4.3.4. We expect this comparison to be more favorable for EGA as the data dimension increases (see Table 2).

6. Application: Learning nonlinear viscoelasticity

In this section, we present a numerical study on the design of uniaxial tests to determine the constitutive parameters of the anisotropic finite-strain viscoelasticity model in Section 2.3. The goal of this study is to demonstrate that the results in Section 5 for the linear model can be generalized to more complicated constitutive models with large deformation. In Section 6.1 we specify the settings of the design problem. Next, we present the results of ESFIM-optimized design optimization in Section 6.2. Compared with the results in Section 5, we additionally visualize and quantitatively interpret designs that are optimized with regularization of the imposed strain rate, yielding numerically identified smooth loading protocols that balance utility maximization and design feasibility. The implementation details are in Appendix C.

6.1. Settings

Experimental design. The uniaxial test has an almost identical setup to that in Section 5.1. Here, however, we consider a larger maximal imposed strain of 50%: $z_l \in [0, 1]^{10}$. \diamond

Reparameterization and the prior distribution. We consider a reparameterization of the nonlinear viscoelastic constitutive model in Section 2.3 with two viscous branches and 15 constitutive parameters. A two-step procedure, as in Section 5.1, is used. This yields a standard normal prior, and the physical parameters are uniformly distributed and defined via a transformation of the prior in Eq. (22).

We describe the physical parameters and justify the associated ranges; see a summary in Table 2. First, we set the bound for the equilibrium shear modulus μ_∞ in the logarithmic scale, and the bulk modulus κ is set using a distribution on the initial (small-strain) Poisson ratio ν to enforce near incompressibility:

$$\log \mu_\infty \sim \mathcal{U}(-1.5, -0.5), \quad \nu \sim \mathcal{U}(0.45, 0.49), \quad \kappa = \mu_\infty \frac{2(1 + \nu)}{3(1 - 2\nu)}.$$

The range of the shear modulus is set to keep the overall reaction force of order 1. The equilibrium initial fiber stiffness is defined relative to the shear modulus through a logarithmic ratio ϕ_f , and the stiffening rate is set to a uniform distribution in the logarithmic scale

$$\phi_f \sim \mathcal{U}(0, 3), \quad k_1 := \mu_\infty \exp(\phi_f), \quad \log k_2 \sim \mathcal{U}(-1, 1).$$

The resulting maximum fiber stiffness is around 20 times the matrix stiffness. We set a uniform prior for fiber orientation angle, $\alpha_f \in \mathcal{U}(-\pi/4, \pi/4)$, where the fiber is assumed to be generally aligned with the uniaxial direction. For each viscous branch, we use logarithmic ratios to relate its parameters to the equilibrium parameters:

$$q_{\mu,i}, q_{k_1,i}, q_{k_2,i} \sim \mathcal{U}(-1.5, 0.7), \quad \mu_i = \mu_\infty \exp(q_{\mu,i}), \quad k_{1,i} = k_1 \exp(q_{k_1,i}), \quad k_{2,i} = k_2 \exp(q_{k_2,i}).$$

We assign a uniform prior on the rate sensitivity exponent $m_i \sim \mathcal{U}(1.5, 3)$. Finally, to ensure the identifiability of the viscous branches, the logarithm of relaxation time for each branch, $\tau_i = (\gamma_i \mu_i)^{-1}$, is assigned to a distinct range,

$$\log \tau_1 \sim \mathcal{U}(-3.0, -1.5), \quad \log \tau_2 \sim \mathcal{U}(-1.5, 0).$$

Consequently, γ_i is replaced by $\log \tau_i$ in the set of constitutive parameters. \diamond

Observation operator. We use the same observation operators as in Section 5.1, except that (i) the field of view has a height of 1.2 and a width of 3 to accommodate the larger maximal prescribed strain, leading to an image size of 600×1550 (height \times width). (ii) the averaged speckle radius is 7×10^{-3} . In Fig. 12, we show an example of the simulated data generated from the forward model. \diamond

6.2. Results

We utilize ESFIM for experimental design optimization. The surrogate FIM construction and the design optimization settings are given as follows.

Surrogate FIM. The surrogate is given by a residual network with 6 hidden layers and GELU activation, where each layer has a width of 240. The surrogate is trained on FIMs computed from 16,384 parameter and design-variable samples generated via QMC. The surrogate FIM achieves a mean relative accuracy of over 97% in predicting the log. FIM on a testing dataset of 4096 samples. Moreover, the surrogate achieves a mean relative accuracy of over 99.4% in predicting the approximate information gain in Eq. (16). In the right panel of Fig. 5, we visualize the surrogate FIM sample complexity, i.e., testing error as a function of training sample size. A standard MC rate of convergence at $N_{\text{train}}^{-1/2}$ is observed. \diamond

Table 2

The bounds for the uniform distributions of the physical parameters in the re-parameterized nonlinear viscoelastic constitutive model.

	Physical parameter	Description	Bound
Equilibrium	$\log \mu_\infty$	Log. of eq. shear modulus	(-1.5, -0.5)
	ν	Initial Poisson ratio	(0.45, 0.49)
	ϕ_f	eq. fiber stiffness-to-shear Log. ratio	(0, 3)
	$\log k_2$	Log. of eq. fiber stiffening rate	(-1, 1)
	α_f	Fiber orientation angle	($-\pi/4, \pi/4$)
Non-Equilibrium	$q_{\mu,i}, i = 1, 2$	Viscous-to-eq. shear Log. ratio	(-1.5, 0.7)
	$q_{k_1,i}, i = 1, 2$	Viscous-to-eq. fiber stiffness Log. ratio	(-1.5, 0.7)
	$q_{k_2,i}, i = 1, 2$	Viscous-to-eq. fiber stiffening rate Log. ratio	(-1.5, 0.7)
	$m_i, i = 1, 2$	Rate-sensitivity exponent	(1.5, 3)
	$\log \tau_1$	Log. of relaxation time (fast)	(-3.0, -1.5)
	$\log \tau_2$	Log. of relaxation time (slow)	(-1.5, 0)

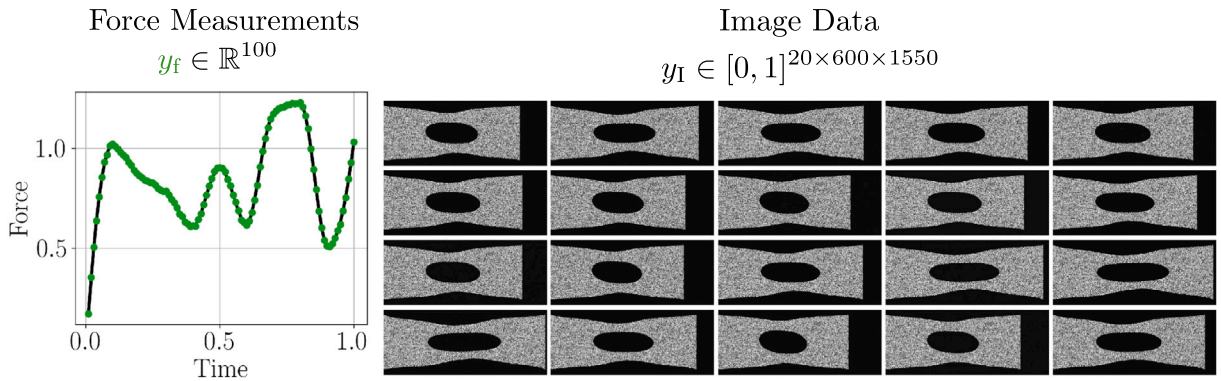


Fig. 12. An example of simulated force and image data for learning nonlinear viscoelasticity. This simulated dataset is generated using the experimental setup in Section 5.1 and Fig. 3, with an order-of-magnitude higher imposed strain. The constitutive parameters are drawn from the prior. Note that the resolution of the image data visualized here is much lower than that of the actual data used in the numerical study; see an image snapshot with high resolution in Fig. C.17.

Design optimization. Two sets of results on design optimization of batch sizes 1–3 are presented: (i) ESFIM maximization with bound constraints and no additional regularization, and (ii) ESFIM maximization with shared geometry, bound constraints, and Tikhonov regularization on the loading path. In particular, we consider the following regularization term for the control points z_i of a single loading path

$$\mathcal{P}(z_i) = \lambda \int_0^1 |\dot{u}_i(s; z_i)|^2 ds + \chi_{[0,1]^{10}}(z_i), \tag{24}$$

where the $\lambda > 0$ is a regularization constant, u_i is the interpolated loading path, χ is the indicator function for enforcing bound constraints, which returns 0 if the control is inside the bounds and infinite if outside. For batched designs, the regularization terms for different experiments are averaged. This regularization term encourages the generated loading protocols to be smooth and physically realizable by eliminating step-like control inputs that require high actuation speeds. \diamond

Now we present and analyze the ESFIM-optimized designs.

6.2.1. Optimized designs

The optimal designs for batch sizes of 1–3 with no regularization are shown in the top panel of Fig. 13. For a single experiment, the design optimization found a loading path that loads to maximum rapidly, holds for 8 control points, and rapidly unloads and holds for the last two control points. As the batch size increases, the additional loading path introduces transitions between maximal strain and no strain. This pattern is similar to the linear model—improving the identifiability of the parameters associated with the instantaneous response becomes more rewarding as the batch size increases. Regarding specimen geometry, a preference for large aspect ratios is observed across all batch sizes. Distinct geometries are introduced as batch size increases: tilted holes with large aspect ratios for batch sizes of 2 and 3, and a small hole with an equal aspect ratio for batch size of 3.

We compare the optimized designs with an ensemble of 160 designs randomly drawn from the uniform distribution within the design variable bounds. The distribution of the EGA values for random designs is shown on the bottom-left panel of Fig. 13, along with the EGA values of the ESFIM-optimized designs. The expected utility improvements in the bottom-right panel of Fig. 13 indicate that the optimized designs increased confidence in parameter identification by approximately 7%–9% compared to random designs of the same batch size, demonstrating ESFIM’s cost-amortization capability and suitability for batched design.

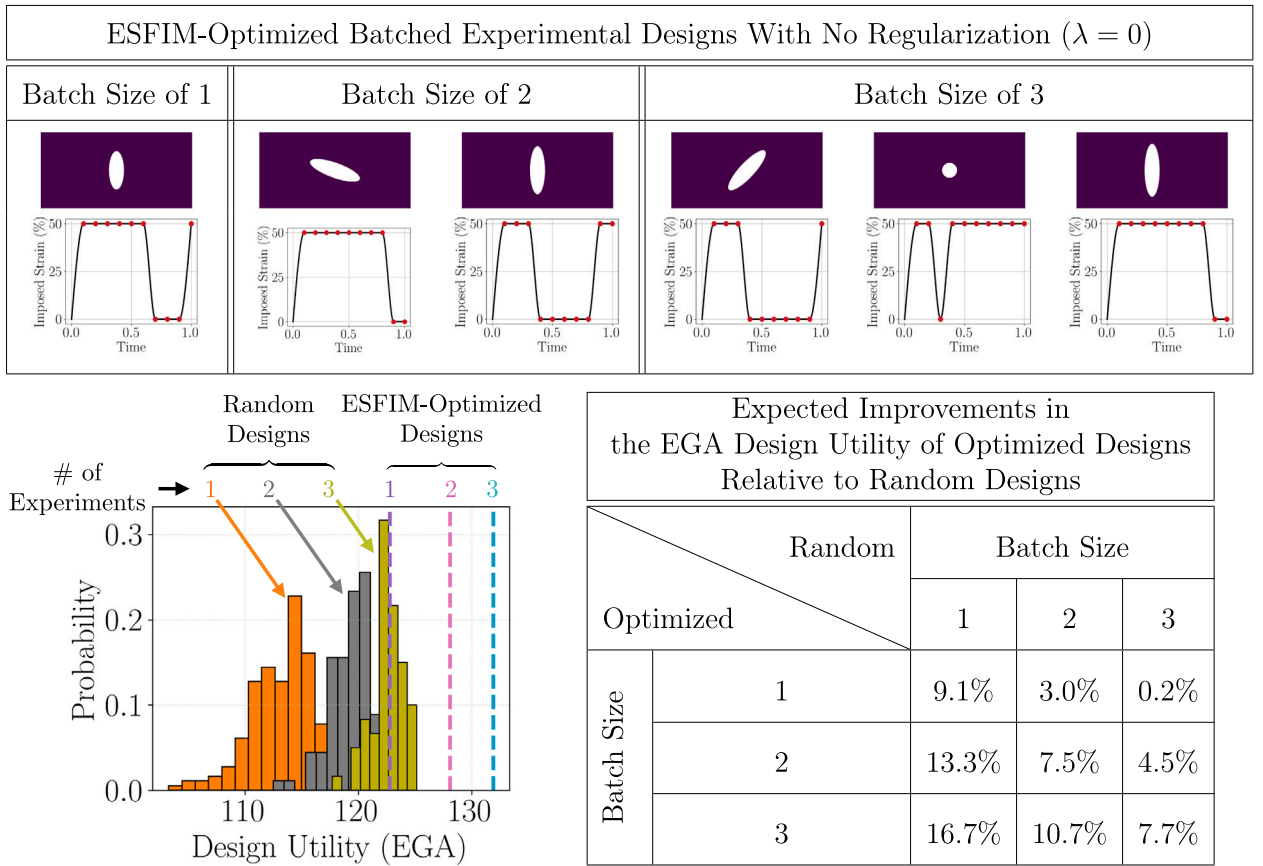


Fig. 13. Visualization and analysis of ESFIM-optimized designs for uniaxial testing of nonlinear viscoelastic materials. (Top) The designs obtained via ESFIM maximization with batch sizes of 1–3. (Bottom Left) The EGA design utility values for random and ESFIM-optimized designs with batch sizes of 1–3. (Bottom Right) The expected improvements in the EGA of the optimized designs with batch sizes of 1–3 relative to random designs with batch sizes of 1–3.

In Fig. 14, we visualize the optimized designs with shared geometry and strain rate regularization. We consider three regularization constants, $\lambda \in \{0.25, 1, 4\}$. The resulting optimized loading path is much smoother compared to those in Fig. 13 with $\lambda = 0$, suggesting that the design optimization finds a balance between ESFIM maximization and regularization minimization. While loading paths with switches between maximal and no-strain are still preferred, intermediate control values are introduced to reduce the imposed strain rate. For $\lambda = 4$ and batch sizes of 2 and 3, one of the loading paths leads to half of the maximum strain, the other loading paths in the same batch load to the maximum strain, indicating a further balance among the multiple regularization terms within the same batch. In general, these results demonstrate the benefit of the ESFIM in numerically discovering non-intuitive experimental designs, with cost amortization across batch sizes and design-optimization settings. In this case, ESFIM enables the study of how regularization constants affect design optimization across varying batch sizes at negligible computational cost beyond constructing the surrogate FIM.

6.2.2. Quantitative interpretations

We provide quantitative interpretations of the ESFIM-optimized design of a single experiment without ($\lambda = 0$, Fig. 13) and with ($\lambda = 0.25$, Fig. 14) strain rate regularization.

In Fig. 15, we present the visualizations and statistics of the 95% credible intervals (CIs) for the posteriors. The credible intervals are again computed using the local Gaussian approximation in Section 4.3.1. In the top panel of Fig. 15, we visualize the distribution of 95% CI sizes across prior-based variation in experimental outcomes. For random designs, the size distribution also accounts for design variations. The optimized designs consistently shift the distribution of 95% CI lengths toward smaller values, with the design obtained without regularization being more effective. In the bottom panel of Fig. 15, we provide the relative expected reduction of the 95% CIs by optimized designs compared to random designs. The design optimization significantly reduces the 95% CI sizes, with average reductions of 44% without regularization and 29% with regularization. We observe that regularization substantially hinders parameter identifiability in the first viscous branch that captures fast stress relaxation.

In Fig. 16, we provide the EGA for inferring four disjoint subsets of parameters. The four subsets corresponding to (i) matrix equilibrium properties, $(\log \mu_\infty, \nu)$, (ii) fiber properties, $(\phi_f, \log k_2, \alpha_f)$, (iii) viscous amplitude, $(q_{\mu,i}, q_{k_1,i}, q_{k_2,i}, i = 1, 2)$, and (iv)

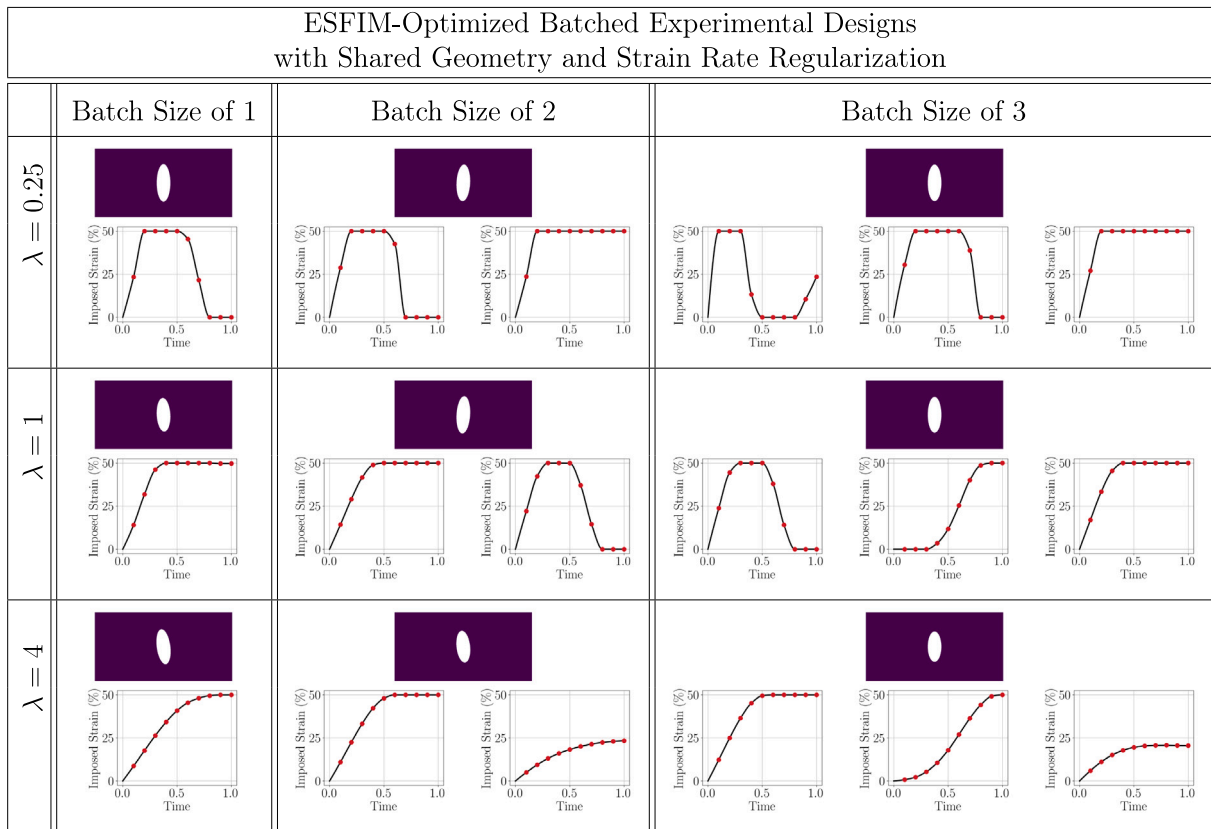


Fig. 14. The experimental designs with batch sizes 1–3 found via ESFIM maximization for uniaxial testing of nonlinear viscoelastic materials. The designs within the same batch share a geometry, and the design optimization is subject to strain-rate regularization in Eq. (24) with three regularization constants $\lambda \in \{0.25, 1, 4\}$.

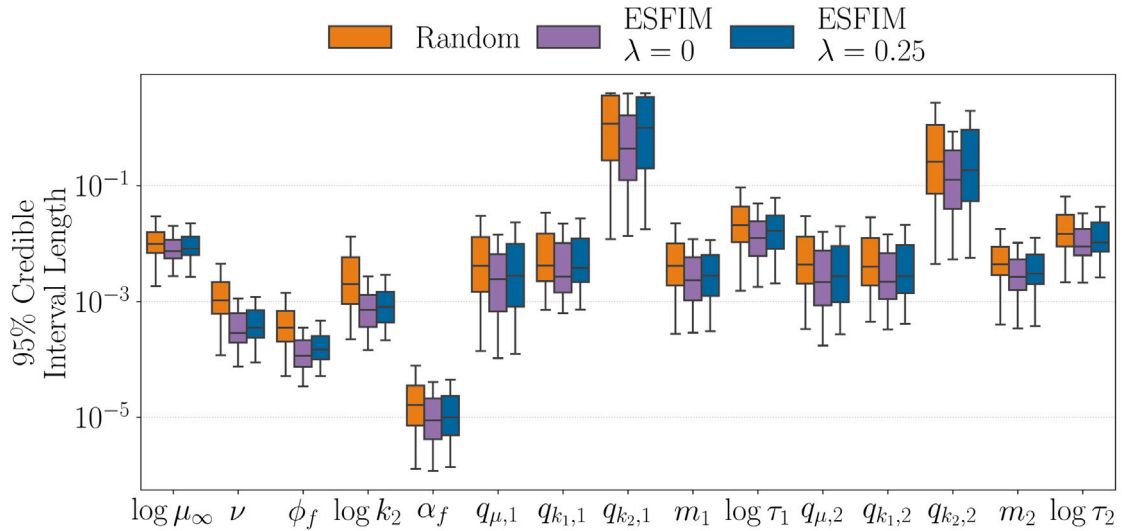
viscous kinetics, $(m_i, \log \tau_i, i = 1, 2)$. The optimized designs consistently improve the identifiability of these parameter subsets compared to random designs. For the design obtained without regularization, the relative utility improvement is greatest for the viscous amplitude (8.3%), whereas the improvement in matrix equilibrium properties is the least significant, yet still 5%. For the design obtained with regularization, the relative utility improvements are smaller, particularly for the matrix equilibrium behavior. The relative improvement of this design is most significant for viscous kinetics, at over 6%, with matrix equilibrium properties the least significant, at 3%.

7. Conclusion and outlook

In this work, we propose a simulation-based framework for quantifying, interpreting, and maximizing the utility of experimental design to reliably learn history-dependent constitutive models. Using a Bayesian formulation of the parameter identification problem, we assess the quality of experimental data by the expected information gain (EIG) to quantify the reduction in parametric uncertainty.

To address the computational challenges associated with EIG estimation, we consider the Bayesian D-optimal design based on a Gaussian approximation of the EIG. The resulting EGA utility not only significantly reduces computational complexity by replacing the nested Monte Carlo sampling of the forward model with a single-layer Monte Carlo sampling of the Fisher information matrix (FIM) but also enables insightful physical interpretations of design utility by analyzing the posterior marginals. Furthermore, we addressed the scalability issues in designing multiple experiments by introducing an estimator based on a surrogate FIM (ESFIM). This approach enables efficient optimization of batched experimental designs by amortizing the surrogate model's training costs across all batch sizes.

We demonstrate the efficacy of this framework through numerical case studies on uniaxial testing for both linear and nonlinear viscoelastic solids. The results indicate that optimized specimen geometries and loading paths yield high-quality force and image data, leading to significantly narrower credible intervals for constitutive parameters than those from random experimental designs. For simple models such as linear viscoelasticity, the framework identified designs that align with our expectations, including tilted elliptical holes and specific loading-holding or unloading-holding sequences that isolate and amplify the signals of anisotropy and

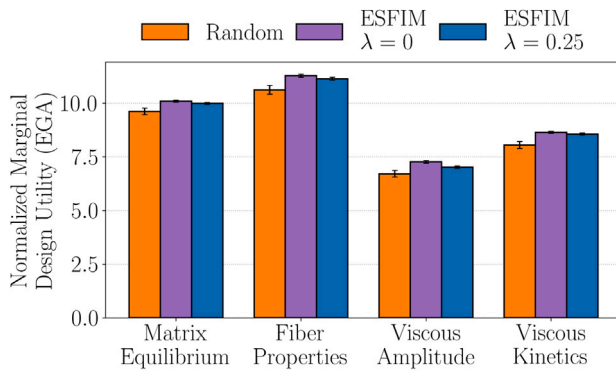


Expected Reduction of the 95% Credible Intervals By Optimized Designs Relative to Random Designs								
	$\log \mu_\infty$	ν	ϕ_f	$\log k_2$	α_f	$q_{\mu,1}$	$q_{k_1,1}$	$q_{k_2,1}$
ESFIM ($\lambda = 0$)	25.5%	76.6%	68.3%	78.6%	40.3%	38.8%	42.8%	32.6%
ESFIM ($\lambda = 0.25$)	16.4%	72.3%	61.3%	76.5%	34.4%	16.1%	10.4%	6.4%
	m_1	$\log \tau_1$	$q_{\mu,2}$	$q_{k_1,2}$	$q_{k_2,2}$	m_2	$\log \tau_2$	Averaged
ESFIM ($\lambda = 0$)	42.3%	41.0%	36.8%	37.8%	33.2%	34.6%	33.2%	44.2%
ESFIM ($\lambda = 0.25$)	14.6%	27.1%	19.9%	19.2%	10.8%	30.0%	22.9%	29.2%

Fig. 15. Visualizations and statistics of the 95% CIs for constitutive parameters of nonlinear viscoelasticity inferred from experiments with random and ESFIM-optimized designs. We consider design optimization without ($\lambda = 0$) and with ($\lambda = 0.25$) strain-rate regularization. (Top) The box plots of the 95% CI sizes. The distributions arise from the prior-based variation in the experimental outcomes. For random designs, the distributions also account for design variations. (Bottom) The expected reduction of the 95% CI sizes by optimized designs relative to random designs. The sizes are estimated for θ , which is transformed as in Section 6.1 from the physical parameters used as labels.

stress relaxation. We also show that the framework can identify non-intuitive designs, such as those for nonlinear viscoelasticity with strain-rate regularization. Overall, this work provides a systematic tool for advanced material testing, enabling the automated discovery of experimental protocols that enrich the information content of experimental data and support reliable learning of constitutive models from designed experiments.

While we demonstrate the potential of the proposed framework to facilitate learning of constitutive models from experiments, several avenues for further improvement and extension remain. Methodologically, the Gaussian approximation to the posterior introduces bias in EIG estimation. Employing a more general non-parametric method for posterior density estimation can enhance the effectiveness of our proposed framework; see, e.g., [86–89,96,97]. There are density estimation methods that can readily incorporate complex, high-dimensional nuisance uncertainty, e.g., uncertain boundary and initial conditions [87], into design utility evaluations to enhance the robustness of parameter identification. In terms of material testing, current numerical studies have primarily focused on uniaxial loading of planar specimens under simplified conditions. Extending this framework to complex three-dimensional geometries and loading scenarios is essential for calibrating more general constitutive models. Such an expansion necessitates integrated studies that simultaneously optimize specimen topology and multi-axial loading protocols, potentially utilizing 3D volumetric data.



Expected Improvement in
the EGA Design Utility
of Optimized Designs
Relative to Random Designs

	ESFIM $\lambda = 0$	ESFIM $\lambda = 0.25$
Matrix Equilibrium	5.0%	3.8%
Fiber Properties	6.3%	4.9%
Viscous Amplitude	8.3%	4.5%
Viscous Kinetics	7.3%	6.3%

Fig. 16. The normalized EGA design utility for learning marginal parameters of nonlinear viscoelasticity grouped by material behaviors. We consider design optimization without ($\lambda = 0$) and with ($\lambda = 0.25$) strain-rate regularization. (Left) The bar plots of the utility values of random and ESFIM-optimized designs for learning matrix equilibrium, fiber, viscous amplitude, and viscous kinetics properties. The error bars for the random designs represent the standard deviations of the utility with respect to the design variables. In contrast, the error bars for the optimized designs represent the standard errors of EGA due to finite parameter sample sizes. (Right) The expected improvement in the utility of optimized design relative to random designs for learning different properties.

CRedit authorship contribution statement

Kaushik Bhattacharya: Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization. **Lianghao Cao:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Andrew Stuart:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work is supported by the ONR MURI, USA on Data-Driven Closure Relations N00014-23-1-2654. AMS is also supported by a Department of Defense Vannevar Bush Faculty Fellowship, USA award number N00014-22-1-2790.

Appendix A. The image observation model

In this appendix, we detail the observation model used for the image data.

A.1. Image deformation via splatting

The image formation operator \mathcal{W} returns the predicted deformed image $I_{\text{pred}} \in [0, 1]^{n \times w \times h}$ given the reference image $I_0 \in [0, 1]^{n \times w \times h}$ and the 2D displacement field $u : \partial\Omega_I \rightarrow \mathbb{R}^2$. Note that here we take u directly as the surface displacement at time t , i.e., $u(\cdot, t)|_{\partial\Omega_I}$, for notational convenience. This operator is implemented using a pushforward algorithm that maps each pixel's intensity in the reference configuration to its new location in the deformed configuration.

Let $\Lambda \subset \mathbb{Z}^2$ denote the pixel indices and $\mathcal{D} \subset \mathbb{R}^2$ denote the physical coordinates of pixels, with $\phi : \Lambda \rightarrow \mathcal{D}$ denoting the index to coordinate map. For each source pixel $p_s \in \Lambda$ of the reference image located in the material domain, i.e., $\phi(p_s) \in \mathcal{D} \cap \partial\Omega_I$, its displacement value $(u \circ \phi)(p_s)$ yields a destination coordinate $x_d = \phi(p_s) + (u \circ \phi)(p_s)$. Since typically $x_d \notin \mathcal{D}$, the source intensity at p_s , denoted by $[I_0]_{p_s}$, is distributed to each of the target pixels located near x_d using a Gaussian kernel $K : \mathbb{R} \rightarrow \mathbb{R}$.

We define an accumulated weight $W \in \mathbb{R}^{n \times w \times h}$ and an accumulated intensity $V \in \mathbb{R}^{n \times w \times h}$. For a target pixel coordinate $p \in \Lambda$, they are computed by summing contributions from all source pixels:

$$[W]_p = \sum_{p_s \in \Lambda} [M_0]_{p_s} K(\|\phi(p) - \phi(p_s) - (u \circ \phi)(p_s)\|; \sigma_s), \quad (\text{A.1})$$

$$[V]_p = \sum_{p_s \in \Lambda} [M_0]_{p_s} [I_0]_{p_s} K(\|\phi(p) - \phi(p_s) - (u \circ \phi)(p_s)\|; \sigma_s), \tag{A.2}$$

where $M_0 \in \{0, 1\}^{n_W \times n_H}$ indicates the pixels of the reference image in the material domain, ensuring that the weights and intensities do not include the image intensities of the background. We use a kernel defined as $K(r; \sigma_s) = \exp(-r^2/2\sigma_s^2)$, where σ_s is the diagonal length of a single pixel. In practice, the kernel is localized via truncation and simplified to a weighted average over a 3×3 pixel window for small deformations and to a 5×5 or larger window for large deformations. The predicted deformed image is obtained by normalizing the accumulated intensity by the accumulated weights:

$$[I_{\text{pred}}]_p = \begin{cases} \frac{[V]_p}{[W]_p} & \text{if } [W]_p > \epsilon, \\ 0 & \text{otherwise,} \end{cases} \tag{A.3}$$

where ϵ is a small threshold to avoid division by zero in regions where no material is present (i.e., the background). This push-forward approach, known as splatting in computer graphics [98], reconstructs the deformed image via kernel regression [99,100], distinct from the pullback interpolation typically used in DIC [46].

A.2. A heteroscedastic noise model with soft masking

We consider an additive noise model for the deformed image, with soft masking to eliminate noise in background pixels. This masking function $\mathcal{A} : [0, 1] \rightarrow [0, 1]$ generates a near-binary mask indicating the region occupied by the deformed material. The mask is computed using a sigmoid function:

$$\mathcal{A}(x) = (1 + \exp(-\beta(x - \gamma)))^{-1}, \tag{A.4}$$

where γ represents a normalized intensity threshold, and β controls the steepness of the transition between the material and the background. We set $\gamma = 0.05$ and $\beta = 10$. The image data y is modeled as the masked prediction perturbed by additive white Gaussian noise η :

$$y = \mathcal{A}(I_{\text{pred}}) \odot (I_{\text{pred}} + \sigma_\eta \eta), \quad [\eta]_{ij} \sim \mathcal{N}(0, 1). \tag{A.5}$$

where \mathcal{A} is applied pixel-wise and \odot is pixel-wise multiplication. The corresponding image misfit function (negative log-likelihood) is given by:

$$\Phi(I_{\text{pred}}) = \frac{1}{2\sigma_\eta^2} \|y \oslash \mathcal{A}(I_{\text{pred}}) - I_{\text{pred}}\|_F^2 - \sum_{p \in \Lambda} \log \mathcal{A}([I_{\text{pred}}]_p), \tag{A.6}$$

where \oslash denotes pixel-wise division.

We contrast our heteroscedastic noise model with two alternatives: hard thresholding and homoscedastic modeling. Hard thresholding discards background pixels in the predicted image, introducing a non-differentiable cutoff that reduces information relevant to boundary features. Homoscedastic modeling assumes that background and material pixels have the same noise level and share the same residual weighting constant. This is an unnatural assumption for the probabilistic modeling of masked data, where background pixels are obtained from post-processing and are noise-free. Our approach resolves these issues by deriving a probabilistic model for masked data. This naturally leads to a negative log-likelihood that strongly penalizes mismatches at material boundaries.

Appendix B. Proof of the EIG equivalence lemma

We consider two factorizations of the joint density of the parameters and data:

$$\pi(\theta, y) = \pi_{\theta|Y(z)}(\theta | y) \pi_{Y(z)}(y) \tag{B.1a}$$

$$= \pi_{Y(z)|\theta}(y | \theta) \pi_\theta(\theta). \tag{B.1b}$$

Expanding the entropy terms in the EIG and applying these factorizations, we have:

$$\begin{aligned} \text{EIG}(z) &= - \int \pi_\theta(\theta) \log \pi_\theta(\theta) d\theta + \iint \pi(\theta, y) \log \pi_{\theta|Y(z)}(\theta | y) d\theta dy && \text{(Apply (B.1a))} \\ &= \iint \pi(\theta, y) \log \frac{\pi_{\theta|Y(z)}(\theta | y)}{\pi_\theta(\theta)} d\theta dy && \left(\pi_\theta = \int \pi(\cdot, y) dy \right) \\ &= \iint \pi(\theta, y) \log \frac{\pi_{Y(z)|\theta}(y | \theta)}{\pi_{Y(z)}(y)} d\theta dy. && \text{(Apply (B.1))} \end{aligned}$$

By factoring the joint density through the posterior, we recover the first form:

$$\begin{aligned} \text{EIG}(z) &= \int \left(\int \log \frac{\pi_{\theta|Y(z)}(\theta | y)}{\pi_\theta(\theta)} \pi_{\theta|Y(z)}(\theta | y) d\theta \right) \pi_{Y(z)}(y) dy && \text{(Apply (B.1a))} \\ &= \mathbb{E}_{y \sim \pi_{Y(z)}} [D_{\text{KL}}(\pi_{\theta|Y(z)}(\cdot | y) \parallel \pi_\theta)]. \end{aligned}$$




Fig. C.17. High resolution snapshots of image data. (Left) The last snapshot in Fig. 4 for uniaxial testing of linear viscoelastic materials. (Right) The last snapshot in Fig. 12 for uniaxial testing of nonlinear viscoelastic materials.

Similarly, by factoring the joint density through the likelihood, we recover the second form:

$$\begin{aligned} \text{EIG}(z) &= \int \left(\int \log \frac{\pi_{Y(z)|\Theta}(y|\theta)}{\pi_{Y(z)}(y)} \pi_{Y(z)|\Theta}(y|\theta) dy \right) \pi_{\Theta}(\theta) d\theta && \text{(Apply (B.1b))} \\ &= \mathbb{E}_{\theta \sim \pi_{\Theta}} [D_{\text{KL}}(\pi_{Y(z)|\Theta}(\cdot|\theta) \parallel \pi_{Y(z)})]. \end{aligned}$$

Appendix C. Numerical implementation details

C.1. Numerical solver

The balance equation and its sensitivity equations are solved using the finite element method implemented with the FEniCS [101, 102] library, with inverse-problem functionality (including the prior, the likelihood, and the symbolic derivation of the sensitivity equations) provided by the hippylib [103] and Geometric MCMC [43] library. In particular, we use linear triangular elements for the displacement field and assign internal variable values at quadrature points, with approximately 17,000 and 30,000 degrees of freedom for the linear and nonlinear models, respectively. In particular, the mesh is refined by a factor of 2.5 and 4 near the elliptical holes for the linear and nonlinear models, respectively. The implicit Euler scheme is used for time evolution. A sparse direct solver is used to solve systems of equations. The Newton method with backtracking line search is used for nonlinear solves. We use 100 time steps for the linear model and 200 time steps for the nonlinear model with adaptive refinement to 300 or more when the Newton method fails to converge. The linear model takes, on average, 10 s to solve, while the nonlinear model takes, on average, 3.2 min. Evaluating the FIM takes, on average, 3.5 min for linear models and 26 min for nonlinear models. The timings are recorded on an Intel Cascadelake 8276 2.1 GHz CPU.

C.2. Image data generation

The synthetic speckle pattern is generated in two stages. First, speckle centers are sampled. To obtain a spatially uniform yet non-periodic distribution, we use a Poisson disk sampling algorithm based on Bridson's method [104], which enforces a minimum distance between any two candidate centers to avoid clustering. In the second stage, we adjust the pattern to match a prescribed area coverage by estimating the number of speckles needed to cover the desired fraction of the domain and randomly subsampling that number of centers. Each retained center is then assigned a radius drawn from a normal distribution around the base radius. The resulting pattern consists of circular speckles with slightly varying size and controlled overall density, providing a realistic, high-contrast pattern.

The image observation model is implemented using JAX [105] to enable automatic differentiation of the image formation operator \mathcal{W} defined in Appendix A. It takes approximately 7 s and 22 s to create the full image data for linear and nonlinear models, respectively. The timings are recorded on an Intel Cascadelake 8276 2.1 GHz CPU. We visualize two high-resolution snapshots of image data in Fig. C.17.

C.3. Principal-stress-based stress-state entropy

We consider the principal-stress-based stress-state entropy following [22,23]. For a given design z , this quantity is evaluated using a fixed set of $N_p = 300$ prior samples $\{\theta^{(i)}\}$.

For each prior sample $\theta^{(i)}$ and design z , we solve the forward model and evaluate the stress tensor at the times where the image snapshots are taken. At each spatial quadrature point and time point, the stress tensor is mapped to its principal stresses, yielding points in a two-dimensional principal-stress space, which produces a stress-state distribution associated with $(\theta^{(i)}, z)$. To quantify the diversity of the stress-state distribution, we discretize the principal-stress space using a rectangular grid shared across all designs

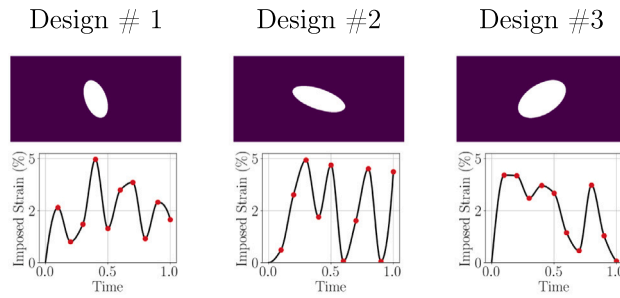


Fig. C.18. Randomly sampled designs used for NMCE evaluations.

and prior samples. Let $p_j(\theta^{(i)}, z)$ denote the empirical probability associated with bin j . The corresponding stress-state entropy is given by

$$H_{ss}(\theta^{(i)}, z) = - \sum_{j: p_j(\theta^{(i)}, z) > 0} p_j(\theta^{(i)}, z) \log p_j(\theta^{(i)}, z).$$

We then define the stress-state entropy of the design z by averaging over the prior samples:

$$\overline{H_{ss}}(z) = \frac{1}{N_p} \sum_{i=1}^{N_p} H_{ss}(\theta^{(i)}, z).$$

This quantity is used to produce Fig. 9.

C.4. Nested Monte Carlo estimator evaluation

To quantify the intractability of the NMC estimator, we estimate the number of inner samples N_{in}^* required to achieve a minimal effective sample size of $ESS \approx 2$. Under a log-normal approximation for the importance weights, one has $ESS \approx N_{in} \exp(-V)$, where $V = \text{Var}[\log w]$ denotes the variance of the log-weights across inner samples. Solving for the sample size yields $N_{in}^* = 2 \exp(V)$, or equivalently $\log_{10} N_{in}^* = (\ln 2 + V) / \ln 10$.

We observe a median log-weight variance of $V \approx 5.9 \times 10^7$, driven almost entirely by the likelihood term rather than the prior-to-proposal reweighting. Substituting V , we obtain $\log_{10} N_{in}^* \approx 2.6 \times 10^7$, i.e., $N_{in}^* \approx 10^{2.6 \times 10^7}$. This sample requirement, which far exceeds any conceivable computational budget, demonstrates that the NMC estimator is fundamentally intractable for this problem, a consequence of the likelihood's extreme sensitivity to perturbations in the model parameters.

We note that the estimate $N_{in}^* \approx 10^{2.6 \times 10^7}$ should be interpreted as an order-of-magnitude diagnostic rather than a precise prediction. The log-normal approximation $ESS \approx N_{in} \exp(-V)$ is exact only when the log-weights are Gaussian. Moreover, since we observed $ESS = 1$ with $N_{in} = 16,384$, the empirical variance V may over- or under-estimate the true log-weight variance. However, because $\log_{10} N_{in}^*$ is proportional to V , even orders-of-magnitude reduction in V (e.g., a better proposal distribution) would only rescale the exponent of N_{in}^* . For instance, a 10^6 times reduction of the empirical variance V would still yield $N_{in}^* \approx 10^{26}$, which remains intractable.

In Fig. C.18, we visualize the randomly sampled designs for NMCE in Section 5.4.

C.5. Surrogate training

Additional details of the surrogate training are provided as follows. The mean relative error in predicting the log-FIM is used as the loss function, where the Frobenius norm for matrices is employed. We use an AdamW optimizer with an initial learning rate of 10^{-3} and a weight decay of 10^{-3} . We use cosine learning rate annealing to decay the learning rate to 10^{-5} at the end of training. We consider 250 total epochs with a batch size of 32. A validation dataset of size 4096 is used to tune the architecture parameters and select the best-performing model during training. The surrogate training is implemented in PyTorch [106].

C.6. Design utility maximization

The EGA maximization is performed using Bayesian optimization implemented through the software library in [107]. The design space is first normalized to the unit hypercube. We employ a Gaussian process surrogate of the EGA utility function with a Matérn kernel ($\nu = 5/2$) and Automatic Relevance Determination [108]:

$$K(z, z') = \sigma_f^2 \left(1 + \sqrt{5}r + \frac{5}{3}r^2 \right) \exp(-\sqrt{5}r), \quad r = \sqrt{\sum_{j=1}^{d_z} \frac{(|z_j| - |z'_j|)^2}{|l_j|^2}}, \quad (C.1)$$

where l_j are the length-scale hyperparameters optimized for each design dimension and σ_f^2 is the signal variance. The search strategy proceeds in two phases: (i) an initialization phase with $N_{\text{init}} = 64$ design samples drawn via QMC, and (ii) an adaptive phase of $N_{\text{iter}} = 268$ iterations. In the adaptive phase, new designs are selected by maximizing the expected improvement acquisition function. To balance exploration and exploitation, the exploration parameter is set to 0.05 and decays by 5% per iteration. The sample-averaged approximation in Eq. (18) is parallelized, where the FIM evaluations at 128 parameter samples occur concurrently. Instead of iid parameter samples in Eq. (18), we consider correlated samples drawn via QMC.

The ESFIM maximization is performed by first evaluating the objective over a large number of randomly sampled design variables (e.g., 10^6), which can be done efficiently on a GPU. Then, a smaller pool of 256 samples with high ESFIM values is selected as initial guesses for ESFIM maximization using the L-BFGS-B algorithm. The final optimal design is selected as the one with the maximum ESFIM across all optimization runs.

Data availability

We have shared code through link to a public Github repository.

References

- [1] A. Kirsch, An introduction to the mathematical theory of inverse problems, third ed., Applied Mathematical Sciences, vol. 120, Springer, Cham, 2021, <http://dx.doi.org/10.1007/978-3-030-63343-1>.
- [2] O. Ghattas, K. Willcox, Learning physics-based models from data: Perspectives from inverse problems and model reduction, *Acta Numer.* 30 (2021) 445–554, <http://dx.doi.org/10.1017/S0962492921000064>.
- [3] A.M. Stuart, Inverse problems: A Bayesian perspective, 2010/05/10, *Acta Numer.* 19 (2010) 451–559, <http://dx.doi.org/10.1017/S0962492910000061>.
- [4] K. Chaloner, I. Verdinelli, Bayesian Experimental Design: A Review, *Statist. Sci.* 10 (3) (1995) 273–304, <http://dx.doi.org/10.1214/ss/1177009939>.
- [5] X. Huan, J. Jagalur, Y. Marzouk, Optimal experimental design: Formulations and computations, *Acta Numer.* 33 (2024) 715–840, <http://dx.doi.org/10.1017/S0962492924000023>.
- [6] 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 (1) (2021) e12370, <http://dx.doi.org/10.1111/str.12370>.
- [7] S. Avril, M. Bonnet, A.-S. Bretelle, M. Grédiac, F. Hild, P. Ienny, F. Latourte, D. Lemosse, S. Pagano, E. Pagnacco, F. Pierron, Overview of identification methods of mechanical parameters based on full-field measurements, *Exp. Mech.* 48 (4) (2008/08/01) 381–402, <http://dx.doi.org/10.1007/s11340-008-9148-y>.
- [8] B. Liu, E. Ocegueda, M. Trautner, A.M. Stuart, K. Bhattacharya, Learning macroscopic internal variables and history dependence from microscopic models, *J. Mech. Phys. Solids* 178 (2023) 105329, <http://dx.doi.org/10.1016/j.jmps.2023.105329>.
- [9] K. Bhattacharya, L. Cao, G. Stepaniants, A. Stuart, M. Trautner, Learning memory and material dependent constitutive laws, 2025, <http://dx.doi.org/10.48550/arXiv.2502.05463>, ArXiv Preprint.
- [10] J.N. Fuhg, G. Anantha Padmanabha, N. Bouklas, B. Bahmani, W. Sun, N.N. Vlassis, M. Flaschel, P. Carrara, L. De Lorenzis, A review on data-driven constitutive laws for solids, *Arch. Comput. Methods Eng.* 32 (3) (2025) 1841–1883, <http://dx.doi.org/10.1007/s11831-024-10196-2>.
- [11] T. Lookman, P.V. Balachandran, D. Xue, R. Yuan, Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design, *Npj Comput. Mater.* 5 (1) (2019) 21.
- [12] S.S. Fayad, E.M.C. Jones, D.T. Seidl, J. Lambros, Identification uncertainty in inverse material model parameter determination: A sensitivity-based decision process for load path selection, *Strain* 61 (3) (2025) e70007, <http://dx.doi.org/10.1111/str.70007>, e70007 9459982.
- [13] F. Pukelsheim, Optimal design of experiments, *Classics in Applied Mathematics*, vol. 50, Society for Industrial and Applied Mathematics, 2006.
- [14] D.V. Lindley, On a measure of the information provided by an experiment, *Ann. Math. Stat.* 27 (4) (1956) 986–1005, <http://dx.doi.org/10.1214/aoms/1177728069>.
- [15] A. Alexanderian, Optimal experimental design for infinite-dimensional Bayesian inverse problems governed by PDEs: a review, *Inverse Problems* 37 (4) (2021) 043001, <http://dx.doi.org/10.1088/1361-6420/abe10c>.
- [16] T. Rainforth, A. Foster, D.R. Ivanova, F.B. Smith, Modern Bayesian experimental design, *Statist. Sci.* 39 (1) (2024) 100–114, <http://dx.doi.org/10.1214/23-STS915>.
- [17] D. Ricciardi, D. Seidl, B. Lester, A. Jones, E. Jones, Bayesian optimal experimental design for constitutive model calibration, *Int. J. Mech. Sci.* 265 (2024) 108881, <http://dx.doi.org/10.1016/j.ijmecsci.2023.108881>.
- [18] T. Chu, J.B. Estrada, S.H. Bryngelson, Bayesian optimal design accelerates discovery of soft material properties from bubble dynamics, *Comput. Mech.* 76 (2) (2025) 431–447, <http://dx.doi.org/10.1007/s00466-025-02606-4>.
- [19] A.R. Castillo, V.R. Joseph, S.R. Kalidindi, Bayesian sequential design of experiments for extraction of single-crystal material properties from spherical indentation measurements on polycrystalline samples, *JOM* 71 (8) (2019) 2671–2679, <http://dx.doi.org/10.1007/s11837-019-03549-x>.
- [20] X. Huan, Y.M. Marzouk, Sequential Bayesian optimal experimental design via approximate dynamic programming, 2016, <http://dx.doi.org/10.48550/arXiv.1604.08320>, ArXiv Preprint.
- [21] R. Villarreal, N.N. Vlassis, N.N. Phan, T.A. Catanach, R.E. Jones, N.A. Trask, S.L.B. Kramer, W. Sun, Design of experiments for the calibration of history-dependent models via deep reinforcement learning and an enhanced Kalman filter, *Comput. Mech.* 72 (1) (2023) 95–124.
- [22] R.C. Ihuanyi, J. Luo, W. Li, J. Zhu, Seeking the most informative design of test specimens for learning constitutive models, *Extrem. Mech. Lett.* 69 (2024) 102169, <http://dx.doi.org/10.1016/j.eml.2024.102169>.
- [23] R.C. Ihuanyi, W. Li, M.Z. Bazant, J. Zhu, Mechanics informatics: A paradigm for efficiently learning constitutive models, *J. Mech. Phys. Solids* 203 (2025) 106239, <http://dx.doi.org/10.1016/j.jmps.2025.106239>.
- [24] M.B.R. Bertin, F. Hild, S. Roux, Optimization of a cruciform specimen geometry for the identification of constitutive parameters based upon full-field measurements, *Strain* 52 (4) (2016) 307–323, <http://dx.doi.org/10.1111/str.12178>.
- [25] L. Chamoin, C. Jailin, M. Diaz, L. Quesada, Coupling between topology optimization and digital image correlation for the design of specimen dedicated to selected material parameters identification, *Int. J. Solids Struct.* 193–194 (2020) 270–286, <http://dx.doi.org/10.1016/j.ijsolstr.2020.02.032>.
- [26] J. Beck, B.M. Dia, L.F. Espath, Q. Long, R. Tempone, Fast Bayesian experimental design: Laplace-based importance sampling for the expected information gain, *Comput. Methods Appl. Mech. Engrg.* 334 (2018) 523–553, <http://dx.doi.org/10.1016/j.cma.2018.01.053>.
- [27] Q. Long, M. Scavino, R. Tempone, S. Wang, Fast estimation of expected information gains for Bayesian experimental designs based on Laplace approximations, *Comput. Methods Appl. Mech. Engrg.* 259 (2013) 24–39, <http://dx.doi.org/10.1016/j.cma.2013.02.017>.

- [28] T. Goda, T. Hironaka, T. Iwamoto, Multilevel Monte Carlo estimation of expected information gains, *Stoch. Anal. Appl.* 38 (4) (2020) 581–600, <http://dx.doi.org/10.1080/07362994.2019.1705168>.
- [29] K. Wu, P. Chen, O. Ghattas, A fast and scalable computational framework for large-scale high-dimensional Bayesian optimal experimental design, *SIAM/ASA J. Uncertain. Quantif.* 11 (1) (2023) 235–261, <http://dx.doi.org/10.1137/21M1466499>.
- [30] A.M. Overstall, J.M. McGree, C.C. Drovandi, An approach for finding fully Bayesian optimal designs using normal-based approximations to loss functions, *Stat. Comput.* 28 (2) (2018/03/01) 343–358, <http://dx.doi.org/10.1007/s11222-017-9734-x>.
- [31] Y. Englezou, T.W. Waite, D.C. Woods, Approximate Laplace importance sampling for the estimation of expected Shannon information gain in high-dimensional Bayesian design for nonlinear models, *Stat. Comput.* 32 (5) (2022/09/30) 82, <http://dx.doi.org/10.1007/s11222-022-10159-2>.
- [32] K.J. Ryan, Estimating expected information gains for experimental designs with application to the random fatigue-limit model, *J. Comput. Graph. Statist.* 12 (3) (2003) 585–603, <http://dx.doi.org/10.1198/1061860032012>.
- [33] T. Rainforth, R. Cornish, H. Yang, A. Warrington, F. Wood, On nesting Monte Carlo estimators, in: J. Dy, A. Krause (Eds.), *Proceedings of the 35th International Conference on Machine Learning*, in: *Proceedings of Machine Learning Research*, 80, PMLR, 2018, pp. 4267–4276, URL <https://proceedings.mlr.press/v80/rainforth18a.html>.
- [34] Q. Long, M. Scavino, R. Tempone, S. Wang, A Laplace method for under-determined Bayesian optimal experimental designs, *Comput. Methods Appl. Mech. Engrg.* 285 (2015) 849–876, <http://dx.doi.org/10.1016/j.cma.2014.12.008>.
- [35] P. Müller, G. Parmigiani, Optimal design via curve fitting of Monte Carlo experiments, *J. Amer. Statist. Assoc.* 90 (432) (1995) 1322–1330.
- [36] A.M. Overstall, D.C. Woods, Bayesian design of experiments using approximate coordinate exchange, *Technometrics* 59 (4) (2017) 458–470, <http://dx.doi.org/10.1080/00401706.2016.1251495>.
- [37] X. Huan, Y.M. Marzouk, Simulation-based optimal Bayesian experimental design for nonlinear systems, *J. Comput. Phys.* 232 (1) (2013) 288–317, <http://dx.doi.org/10.1016/j.jcp.2012.08.013>.
- [38] P. Pandita, P. Tsilifis, N.M. Awalgaonkar, I. Bilonis, J. Panchal, Surrogate-based sequential Bayesian experimental design using non-stationary Gaussian processes, *Comput. Methods Appl. Mech. Engrg.* 385 (2021) 114007, <http://dx.doi.org/10.1016/j.cma.2021.114007>.
- [39] K. Wu, T. O’Leary-Roseberry, P. Chen, O. Ghattas, Large-scale Bayesian optimal experimental design with derivative-informed projected neural network, *J. Sci. Comput.* 95 (1) (2023) <http://dx.doi.org/10.1007/s10915-023-02145-1>.
- [40] J. Go, P. Chen, Accurate, scalable, and efficient Bayesian optimal experimental design with derivative-informed neural operators, *Comput. Methods Appl. Mech. Engrg.* 438 (2025) 117845, <http://dx.doi.org/10.1016/j.cma.2025.117845>.
- [41] S. Lan, T. Bui-Thanh, M. Christie, M. Girolami, Emulation of higher-order tensors in manifold Monte Carlo methods for Bayesian inverse problems, *J. Comput. Phys.* 308 (2016) 81–101, <http://dx.doi.org/10.1016/j.jcp.2015.12.032>.
- [42] T. O’Leary-Roseberry, P. Chen, U. Villa, O. Ghattas, Derivative-informed neural operator: An efficient framework for high-dimensional parametric derivative learning, *J. Comput. Phys.* 496 (2024) 112555, <http://dx.doi.org/10.1016/j.jcp.2023.112555>.
- [43] L. Cao, T. O’Leary-Roseberry, O. Ghattas, Derivative-informed neural operator acceleration of geometric MCMC for infinite-dimensional Bayesian inverse problems, *J. Mach. Learn. Res.* 26 (78) (2025) 1–68, URL <http://jmlr.org/papers/v26/24-0745.html>.
- [44] T.C. Chu, W.F. Ranson, M.A. Sutton, Applications of digital-image-correlation techniques to experimental mechanics, *Exp. Mech.* 25 (3) (1985) 232–244, <http://dx.doi.org/10.1007/BF02325092>.
- [45] F. Hild, S. Roux, Digital image correlation: from displacement measurement to identification of elastic properties – a review, *Strain* 42 (2) (2006) 69–80, <http://dx.doi.org/10.1111/j.1475-1305.2006.00258.x>.
- [46] M.A. Sutton, J.-J. Orteu, H.W. Schreier, *Image correlation for shape, motion and deformation measurements: Basic concepts, theory and applications*, Springer, New York, NY, 2009, <http://dx.doi.org/10.1007/978-0-387-78747-3>.
- [47] P.F. Luo, Y.J. Chao, M.A. Sutton, W.H. Peters, Accurate measurement of three-dimensional deformations in deformable and rigid bodies using computer vision, *Exp. Mech.* 33 (2) (1993) 123–132.
- [48] B. Pan, K. Qian, H. Xie, A. Asundi, Two-dimensional digital image correlation for in-plane displacement and strain measurement: A review, *Meas. Sci. Technol.* 20 (6) (2009) 062001, <http://dx.doi.org/10.1088/0957-0233/20/6/062001>.
- [49] B.K. Bay, T.S. Smith, D.P. Fyhrie, M. Saad, Digital volume correlation: Three-dimensional strain mapping using X-ray tomography, *Exp. Mech.* 39 (3) (1999) 217–226, <http://dx.doi.org/10.1007/BF02323555>.
- [50] A. Buljac, C. Jailin, A. Mendoza, J. Neggens, T. Taillandier-Thomas, A. Bouterf, B. Smaniotto, F. Hild, S. Roux, Digital volume correlation: Review of progress and challenges, *Exp. Mech.* 58 (5) (2018) 661–708, <http://dx.doi.org/10.1007/s11340-018-0390-7>.
- [51] J. Yang, L. Hazlett, A.K. Landauer, C. Franck, Augmented Lagrangian digital volume correlation (ALDVC), *Exp. Mech.* 60 (9) (2020) 1205–1223, <http://dx.doi.org/10.1007/s11340-020-00607-3>.
- [52] M. Gonçalves, A. Andrade-Campos, B. Barroqueiro, On the design of mechanical heterogeneous specimens using multilevel topology optimization, *Adv. Eng. Softw.* 175 (2023) 103314, <http://dx.doi.org/10.1016/j.advengsoft.2022.103314>.
- [53] S. Ghoul, M. Flaschel, S. Kumar, L. De Lorenzis, A topology optimisation framework to design test specimens for one-shot identification or discovery of material models, *J. Mech. Phys. Solids* 203 (2025) 106210, <http://dx.doi.org/10.1016/j.jmps.2025.106210>.
- [54] K.T. Kavanagh, R.W. Clough, Finite element applications in the characterization of elastic solids, *Int. J. Solids Struct.* 7 (1) (1971) 11–23, [http://dx.doi.org/10.1016/0020-7683\(71\)90015-1](http://dx.doi.org/10.1016/0020-7683(71)90015-1).
- [55] B. Chen, B. Starman, M. Halilović, L.A. Berglund, S. Coppieters, Finite element model updating for material model calibration: A review and guide to practice, *Arch. Comput. Methods Eng.* 32 (4) (2025) 2035–2112, <http://dx.doi.org/10.1007/s11831-024-10200-9>.
- [56] M. Grédiac, F. Pierron, S. Avril, E. Toussaint, The virtual fields method for extracting constitutive parameters from full-field measurements: A review, *Strain* 42 (4) (2006) 233–253, <http://dx.doi.org/10.1111/j.1475-1305.2006.tb01504.x>.
- [57] F. Pierron, M. Grédiac, *The virtual fields method: Extracting constitutive mechanical parameters from full-field deformation measurements*, Springer, New York, NY, 2012, <http://dx.doi.org/10.1007/978-1-4614-1824-5>.
- [58] A. Akerson, A. Rajan, K. Bhattacharya, Learning constitutive relations from experiments: 1. PDE constrained optimization, *J. Mech. Phys. Solids* 201 (2025) 106128, <http://dx.doi.org/10.1016/j.jmps.2025.106128>.
- [59] A.A. Oberai, N.H. Gokhale, G.R. Feijóo, Solution of inverse problems in elasticity imaging using the adjoint method, *Inverse Problems* 19 (2) (2003) 297, <http://dx.doi.org/10.1088/0266-5611/19/2/304>.
- [60] G. Besnard, F. Hild, S. Roux, “Finite-element” displacement fields analysis from digital images: Application to Portevin–Le Châtelier bands, *Exp. Mech.* 46 (6) (2006) 789–803, <http://dx.doi.org/10.1007/s11340-006-9824-8>.
- [61] F. Mathieu, H. Leclerc, F. Hild, S. Roux, Estimation of elastoplastic parameters via weighted FEM and Integrated-DIC, *Exp. Mech.* 55 (1) (2015) 105–119, <http://dx.doi.org/10.1007/s11340-014-9888-9>.
- [62] A. Wihardja, K. Bhattacharya, Constitutive relations from images, *J. Appl. Mech.* 92 (8) (2025) 081009, <http://dx.doi.org/10.1115/1.4068870>.
- [63] J. Kirchhoff, D. Luo, T. O’Leary-Roseberry, O. Ghattas, Inference of heterogeneous material properties via infinite-dimensional integrated DIC, 2024, <http://dx.doi.org/10.48550/arXiv.2408.10217>, ArXiv Preprint.
- [64] A.P. Ruybalid, J.P.M. Hoefnagels, O. van der Sluis, M.G.D. Geers, Comparison of the identification performance of conventional FEM updating and integrated DIC, *Internat. J. Numer. Methods Engrg.* 106 (4) (2016) 298–320, <http://dx.doi.org/10.1002/nme.5127>.
- [65] J.R. Janesick, *Photon Transfer*, SPIE Press, Bellingham, WA, 2007, <http://dx.doi.org/10.1117/3.725073>.

- [66] A. Foi, M. Trimeche, V. Katkovnik, K. Egiazarian, Practical Poissonian-Gaussian noise modeling and fitting for single-image raw-data, *Trans. Img. Proc.* 17 (10) (2008) 1737–1754, <http://dx.doi.org/10.1109/TIP.2008.2001399>.
- [67] J. Chao, E.S. Ward, R.J. Ober, Fisher information theory for parameter estimation in single molecule microscopy: Tutorial, *J. Opt. Soc. Amer. A* 33 (7) (2016) B36–B57, <http://dx.doi.org/10.1364/JOSAA.33.000B36>.
- [68] H.W. Schreier, J.R. Braasch, M.A. Sutton, Systematic errors in digital image correlation caused by intensity interpolation, *Opt. Eng., Bellingham* 39 (11) (2000) 2915–2921, <http://dx.doi.org/10.1117/1.1314593>.
- [69] B. Blaysat, M. Grédiac, F. Sur, Effect of interpolation on noise propagation from images to DIC displacement maps, *Internat. J. Numer. Methods Engrg.* 108 (3) (2016) 213–232, <http://dx.doi.org/10.1002/nme.5212>.
- [70] H. Rappel, L.A.A. Beex, J.S. Hale, L. Noels, S.P.A. Bordas, A tutorial on Bayesian inference to identify material parameters in solid mechanics, *Arch. Comput. Methods Engrg.* 27 (2) (2020) 361–385, <http://dx.doi.org/10.1007/s11831-018-09311-x>.
- [71] M. Bhattacharyya, P. Feissel, A Bayesian inference approach for parametric identification through optimal control method, *Internat. J. Numer. Methods Engrg.* 124 (14) (2023) 3145–3165, <http://dx.doi.org/10.1002/nme.7242>.
- [72] A. Jafari, K. Vlachas, E. Chatzi, J.F. Unger, A Bayesian framework for constitutive model identification via use of full field measurements, with application to heterogeneous materials, *Comput. Methods Appl. Mech. Engrg.* 433 (2025) 117489, <http://dx.doi.org/10.1016/j.cma.2024.117489>.
- [73] A. Touminet, S. Cantournet, V. Fabre, P. Kerfriden, A Bayesian extension to FEMU for identification of spatially varying stochastic elastic properties from digital image and volume correlation measurements, *J. Comput. Phys.* 531 (2025) 113946, <http://dx.doi.org/10.1016/j.jcp.2025.113946>.
- [74] M.A.S. Mahmoud, D. Renner, A. Khosravani, S.R. Kalidindi, Sequential Bayesian inference of the GTN damage model using multimodal experimental data, 2025, <http://dx.doi.org/10.48550/arXiv.2510.01016>, ArXiv Preprint.
- [75] L. Gaynutdinova, O. Rokoš, J. Havelka, I. Pultarová, J. Zeman, Bayesian approach to micromechanical parameter identification using integrated digital image correlation, *Int. J. Solids Struct.* 280 (2023) 112388, <http://dx.doi.org/10.1016/j.ijsolstr.2023.112388>.
- [76] B.D. Coleman, M.E. Gurtin, Thermodynamics with internal state variables, *J. Chem. Phys.* 47 (2) (1967) 597–613, <http://dx.doi.org/10.1063/1.1711937>.
- [77] S. Reese, S. Govindjee, A theory of finite viscoelasticity and numerical aspects, *Int. J. Solids Struct.* 35 (26) (1998) 3455–3482, [http://dx.doi.org/10.1016/S0020-7683\(97\)00217-5](http://dx.doi.org/10.1016/S0020-7683(97)00217-5).
- [78] R. Lakes, *Viscoelastic Materials*, Cambridge University Press, 2009.
- [79] G.A. Holzapfel, T.C. Gasser, A viscoelastic model for fiber-reinforced composites at finite strains: Continuum basis, computational aspects and applications, *Comput. Methods Appl. Mech. Engrg.* 190 (34) (2001) 4379–4403, [http://dx.doi.org/10.1016/S0045-7825\(00\)00323-6](http://dx.doi.org/10.1016/S0045-7825(00)00323-6).
- [80] G.A. Holzapfel, R.W. Ogden, On the tension–compression switch in soft fibrous solids, *Eur. J. Mech. A Solids* 49 (2015) 561–569, <http://dx.doi.org/10.1016/j.euromechsol.2014.09.005>.
- [81] T. Nguyen, R. Jones, B. Boyce, Modeling the anisotropic finite-deformation viscoelastic behavior of soft fiber-reinforced composites, *Int. J. Solids Struct.* 44 (25) (2007) 8366–8389, <http://dx.doi.org/10.1016/j.ijsolstr.2007.06.020>.
- [82] M. Ortiz, L. Stainier, The variational formulation of viscoplastic constitutive updates, *Comput. Methods Appl. Mech. Engrg.* 171 (3) (1999) 419–444, [http://dx.doi.org/10.1016/S0045-7825\(98\)00219-9](http://dx.doi.org/10.1016/S0045-7825(98)00219-9).
- [83] J.C. Simo, T.J.R. Hughes, *Computational Inelasticity, Interdisciplinary Applied Mathematics*, vol. 7, Springer, New York, NY, 1998, <http://dx.doi.org/10.1007/b98904>.
- [84] M.S. Bartlett, Approximate confidence intervals, *Biometrika* 40 (1–2) (1953) 12–19, <http://dx.doi.org/10.1093/biomet/40.1-2.12>.
- [85] D. Barber, F. Agakov, The IM algorithm: A variational approach to information maximization, in: *Proceedings of the 17th International Conference on Neural Information Processing Systems, NIPS '03*, MIT Press, Cambridge, MA, USA, 2003, pp. 201–208.
- [86] A. Foster, M. Jankowiak, E. Bingham, P. Horsfall, Y.W. Teh, T. Rainforth, N. Goodman, Variational Bayesian optimal experimental design, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, 32, Curran Associates, Inc., 2019, URL https://proceedings.neurips.cc/paper_files/paper/2019/file/d55cbf210f175f4a37916eafe6c040d-Paper.pdf.
- [87] R. Baptista, L. Cao, J. Chen, O. Ghattas, F. Li, Y.M. Marzouk, J.T. Oden, Bayesian model calibration for block copolymer self-assembly: Likelihood-free inference and expected information gain computation via measure transport, *J. Comput. Phys.* 503 (2024) 112844, <http://dx.doi.org/10.1016/j.jcp.2024.112844>.
- [88] J. Dong, C. Jacobsen, M. Khaloufi, M. Akram, W. Liu, K. Duraisamy, X. Huan, Variational Bayesian optimal experimental design with normalizing flows, *Comput. Methods Appl. Mech. Engrg.* 433 (2025) 117457, <http://dx.doi.org/10.1016/j.cma.2024.117457>.
- [89] F. Li, R. Baptista, Y. Marzouk, Expected information gain estimation via density approximations: Sample allocation and dimension reduction, 2025, <http://dx.doi.org/10.48550/arXiv.2411.08390>, ArXiv Preprint.
- [90] A.K. Saibaba, A. Alexanderian, I.C.F. Ipsen, Randomized matrix-free trace and log-determinant estimators, *Numer. Math.* 137 (2) (2017) 353–395, <http://dx.doi.org/10.1007/s00211-017-0880-z>.
- [91] A. Alexanderian, N. Petra, G. Stadler, O. Ghattas, A-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems with regularized ℓ_0 -sparsification, *SIAM J. Sci. Comput.* 36 (5) (2014) A2122–A2148, <http://dx.doi.org/10.1137/130933381>.
- [92] A. Alexanderian, A.K. Saibaba, Efficient D-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems, *SIAM J. Sci. Comput.* 40 (5) (2018) A2956–A2985, <http://dx.doi.org/10.1137/17M115712X>.
- [93] R. Nickl, *Bayesian non-linear statistical inverse problems*, Zurich Lectures in Advanced Mathematics, vol. 30, European Mathematical Society, Berlin, Germany, 2023.
- [94] A.B. Owen, Scrambling sobol' and Niederreiter–Xing points, *J. Complexity* 14 (4) (1998) 466–489, <http://dx.doi.org/10.1006/jcom.1998.0487>.
- [95] A. Bartuska, A.G. Carlon, L. Espath, S. Krumscheid, R. Tempone, Double-loop randomized quasi-Monte Carlo estimator for nested integration, 2025, <http://dx.doi.org/10.48550/arXiv.2302.14119>, ArXiv Preprint.
- [96] K. Koval, R. Herzog, R. Scheichl, Tractable optimal experimental design using transport maps*, *Inverse Problems* 40 (12) (2024) 125002, <http://dx.doi.org/10.1088/1361-6420/ad8260>.
- [97] L. Cao, J. Chen, M. Brennan, T. O'Leary-Roseberry, Y. Marzouk, O. Ghattas, Lazydino: Fast, scalable, and efficiently amortized Bayesian inversion via structure-exploiting and surrogate-driven measure transport, *J. Mach. Learn. Res.* 27 (24) (2026) 1–71, URL <http://jmlr.org/papers/v27/25-0858.html>.
- [98] L. Westover, Footprint evaluation for volume rendering, *SIGGRAPH Comput. Graph.* 24 (4) (1990) 367–376, <http://dx.doi.org/10.1145/97880.97919>.
- [99] E.A. Nadaraya, On estimating regression, *Theory Probab. Appl.* 9 (1) (1964) 141–142, <http://dx.doi.org/10.1137/1109020>.
- [100] G.S. Watson, *Smooth regression analysis*, *Sankhyā: Indian J. Stat. Ser. A* (1961-2002) 26 (4) (1964) 359–372.
- [101] A. Logg, K.-A. Mardal, G. Wells (Eds.), Automated solution of differential equations by the finite element method, in: *Lecture Notes in Computational Science and Engineering*, vol. 84, Springer, Berlin, Heidelberg, 2012, <http://dx.doi.org/10.1007/978-3-642-23099-8>.
- [102] M.S.A. s, 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, *Arch. Numer. Softw.* 3 (100) (2015) 9–23, <http://dx.doi.org/10.11588/ans.2015.100.20553>.
- [103] 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 Trans. Math. Software* 47 (2) (2021) <http://dx.doi.org/10.1145/3428447>.
- [104] R. Bridson, Fast Poisson disk sampling in arbitrary dimensions, in: *ACM SIGGRAPH 2007 Sketches*, SIGGRAPH '07, Association for Computing Machinery, New York, NY, USA, 2007, pp. 22–es, <http://dx.doi.org/10.1145/1278780.1278807>.
- [105] J. Bradbury, R. Frostig, P. Hawkins, M.J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, Q. Zhang, JAX: composable transformations of Python+NumPy programs, 2018, URL <http://github.com/jax-ml/jax>.

- [106] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, Pytorch: an imperative style, high-performance deep learning library, in: *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Curran Associates Inc., Red Hook, NY, USA, 2019.
- [107] F. Nogueira, Bayesian Optimization: Open source constrained global optimization tool for Python, 2014–, URL <https://github.com/bayesian-optimization/BayesianOptimization>.
- [108] J. Snoek, H. Larochelle, R.P. Adams, Practical Bayesian optimization of machine learning algorithms, in: F. Pereira, C. Burges, L. Bottou, K. Weinberger (Eds.), *Advances in Neural Information Processing Systems*, vol. 25, Curran Associates, Inc., 2012, URL https://proceedings.neurips.cc/paper_files/paper/2012/file/05311655a15b75fab86956663e1819cd-Paper.pdf.