

# Discovery of Physics and Characterization of Microstructure from Data with Bayesian Hidden Physics Models

Steven Atkinson,<sup>1</sup> Yiming Zhang,<sup>1</sup> Liping Wang<sup>1</sup>

<sup>1</sup>GE Research

1 Research Circle

Niskayuna, New York 12309

{steven.atkinson1, yiming.zhang, wangli}@ge.com

## Abstract

There has been a surge in the interest of using machine learning techniques to assist in the scientific process of formulating knowledge to explain observational data. We demonstrate the use of Bayesian Hidden Physics Models to first uncover the physics governing the propagation of acoustic impulses in metallic specimens using data obtained from a pristine sample. We then use the learned physics to characterize the microstructure of a separate specimen with a surface-breaking crack flaw. Remarkably, we find that the physics learned from the first specimen allows us to understand the backscattering observed in the latter sample, a qualitative feature that is wholly absent from the specimen from which the physics were inferred. The backscattering is explained through inhomogeneities of a latent spatial field that can be recognized as the speed of sound in the media.

## Introduction

There is a workflow found ubiquitously throughout many disciplines at the intersection of science and engineering. It begins when one observes some novel phenomenon or behavior. One then designs a laboratory experiment to collect measurements to observe the phenomenon under a controlled setting. Given these observations, one infers relevant physical laws from those observations. Given this new knowledge, one may proceed with constructing physically-accurate computational models of the discovered physics. With this tool developed, hypothetical simulations may be carried out *in silico* for the purposes of engineering design and analysis at considerably reduced expense.

Perhaps the most challenging element of the pipeline is the discovery of the relevant physics from the laboratory data. In fact, there has been a surge of popularity in recent years applying black-box data-driven modeling techniques which involve fitting some general function model to the observation data without the structure of a physical model (??). While such models tend to be highly general, they struggle to give accurate predictions or possess the sort of generalization that enables significant scientific insight. The former challenge may be somewhat addressed if it is inexpensive to obtain massive amounts of observational data, though convergence

may be unsatisfactorily slow due to the large hypothesis space encoded in typical black-box models. The latter is more challenging, as extrapolation in data-driven models is notoriously suspect. For cases in which high accuracy and extrapolation are required, a structured approach that incorporates learned physics seems to be the most effective approach.

In this paper, we consider the above questions in the context of characterizing metal specimens given measurements of their mechanical response to acoustic impulses with a primary focus on discovery of physics from laboratory data. We consider several of approaches, including end-to-end differentiable physical simulations (Degraeve et al. 2019) and physics-informed neural networks (PINNs) (Raissi, Perdikaris, and Karniadakis 2019), eventually culminating in an application of the recently-developed Bayesian Hidden Physics Model (Atkinson 2020) to learn the physics of the system and apply it in an extrapolatory regime to understand behavior qualitatively different from that exhibited by the data from which the physics were discovered, i.e. backscattering due to a surface-breaking crack. A key goal is to understand this phenomenon by building physical knowledge from limited data and a broad hypothesis space. Our machine learning approach contrasts with traditional approaches which seek to calibrate a low-dimensional parameter set for known physics, which typically require considerable domain expertise. Rather, we would like to be led to the appropriate explanation by the data, as one might need to do in the absence of expert knowledge or in a completely novel setting where such knowledge does not yet exist.

The contributions of this paper are as follows:

- We look at the inverse problem on pristine and cracked specimens through both PINNs and a traditional end-to-end differentiable solver that enforces exact agreement on the physics by design. We see that the two qualitatively disagree on their characterization of the cracked specimen, prompting the hypothesis that the wave equation is deficient and that the soft agreement between the wave speed field and the deflection field in PINNs is hiding some missing physics.
- We use the technique of Bayesian Hidden Physics Models first demonstrated in (Atkinson 2020) to discover a novel differential equation that better describes the data from a

hypothesis space containing the wave equation as a special case.

- We demonstrate that the physics learned from the pristine specimen may be transferred to explain the presence of backscattering in a specimen with a flaw, demonstrating the data efficiency and validity in extrapolation conferred by a physics-informed approach.

The significance of this work is the demonstration that Bayesian Hidden Physics Models may be fruitfully applied to discover physics from real-world data sets, suggesting that the end-to-end scientific workflow described above may be realized.

### Problem statement

Consider a physical system with a scalar spatiotemporal observable in two-dimensional space represented as a function  $u(x, y, t)$ . First, we assume that the evolution of  $u$  is described by a dynamical partial differential equation of the form

$$u_{tt} = \mathcal{F}[u; \theta_{\mathcal{F}}] \text{ in } \Omega_{st} = \Omega_s \times [0, T], \quad (1)$$

where  $\theta_{\mathcal{F}}$  denotes the (potentially infinite-dimensional) parameterization of the unknown operator  $\mathcal{F}$ . The initial-boundary conditions accompanying Eq. (1) are generally not known. In addition to Eq. (1), one also has access to some finite set of observation data on  $u$ , denoted as  $\mathcal{D} = \{x^{(i)}, y^{(i)}, t^{(i)}, \hat{u}^{(i)}\}_{i=1}^n$ . Third, we assume that one has access to a simulation code capable of solving Eq. (1). The simulation code may require the user to specify initial-boundary conditions with parameters  $\{\theta_{IC}, \theta_{BC}\}$ . The simulation code may also approximate  $\theta_{\mathcal{F}}$  with some internal, finite-dimensional representation  $\theta_f$ . In sum, the parameters of the simulation code are denoted  $\theta = (\theta_f, \theta_{IC}, \theta_{BC})$ . Assume that the output of the simulator may be queried as  $\mathbf{u} = (u(x^{(1)}, y^{(1)}, t^{(1)}; \theta), \dots, u(x^{(n)}, y^{(n)}, t^{(n)}; \theta)) \in \mathbb{R}^n$ . The goal of the simulator calibration problem with known physics is to determine optimal values for theta  $\theta$  such that optimally explain  $\mathcal{D}$  in some sense.

In the traditional setting, optimality may be defined in terms of some loss function to be minimized such as mean squared error. However, even in the presence of a large volume of observation data, its coverage in the space relevant to the input of the differential operator may be sparse (Atkinson 2020), yielding many potential parameter sets  $(\theta_f)$  consistent with  $\mathcal{D}$ . Furthermore, measurements in  $\mathcal{D}$  may be noisy. To address these challenges, we will consider a Bayesian approach by defining a prior  $p(\theta_f)$ , a likelihood associated with the observed data  $p(\hat{u}|u; \theta_{\hat{u}})$  with parameters  $\theta_{\hat{u}}$ , and use variational inference (Blei, Kucukelbir, and McAuliffe 2017) to approximate its posterior by maximizing an evidence lower bound,

$$ELBO(\theta; \mathcal{D}) = \langle p(\hat{u}|u(\theta); \theta_{\hat{u}}) \rangle_{q(\theta_f|\theta_{\setminus f})} - \text{KL}(q(\theta_f|\theta_{\setminus f}) || p(\theta_f|\theta_{\setminus f})), \quad (2)$$

where  $\theta_{\setminus f} = (\theta_{IC}, \theta_{BC}, \theta_{\hat{u}})$  includes all parameters over which a point estimate will be taken. Note that while this approach might be extended to  $\theta_{IC}$ ,  $\theta_{BC}$ , and  $\theta_{\hat{u}}$ , in practice the uncertainty associated with these are negligible compared to that on  $\theta_f$  if many data are available (Atkinson 2020).

## Methodology

Here, we review the methodologies used for microstructure characterization. We first discuss the acquisition of the laboratory dataset subject to analysis. We then discuss how characterization may be solved as an inverse problem via gradient-based calibration through a differentiable simulation code or the method of physics-informed neural networks for inverse problems. Finally, we review Bayesian Hidden Physics Models for discovery of nonlinear partial differential operators with uncertainty quantification as a means of solving the inverse problem without knowing *a priori* what the governing physics are.

### Data acquisition

We base our analysis on datasets from two specimens. The first, referred to as the ‘‘pristine’’ specimen, is a nickel polycrystalline alloy. The second, also used in (Shukla et al. 2020) and referred to as the ‘‘cracked’’ specimen, is a National Institute of Standards and Technology (NIST) surface-breaking crack reference standard (RM 8458) in a  $7 \text{ cm} \times 7 \text{ cm} \times 2 \text{ cm}$  block of 7075-T651 Aluminum alloy substrate material. To generate the datasets, a piezo-electric transducer imparts a wave packet with a frequency of 5 MHz into the specimen, causing a wavefield to travel across the surface of the specimen. For the cracked specimen, the wave packet is propagated normal to the crack line. The wavefield is measured via laser doppler vibrometry at a temporal resolution of 50 MHz ( $0.02 \mu\text{s}$  per sample) and spatial resolution of 0.05 mm. For the pristine specimen, the dataset comprises 1024 measurements in time, on a grid of  $400 \times 260$  spatial locations. For the cracked specimen, the dataset comprises 1024 measurements in time, on a grid of  $240 \times 240$  spatial locations.

### Modeling approaches

We now discuss the various modeling approaches we apply to understand the data described above. A key desideratum is that the relationship between the representation of the microstructure in our models and the resultant measurement field (sampled at a finite set of discrete points) be differentiable. Therefore, we opt for a heterogeneous continuum model of the specimen. It is known that, under the assumption that the physics are represented by the (two-dimensional) wave equation,

$$u_{tt}(x, y, t) = v^2(x, y)\Delta u(x, y, t), \quad (3)$$

the presence of microstructure flaws may be modeled as a locally-depressed value of  $v$  (Seidl 2018). Our construction reflects the physical insight that flaws in metallic media may be captured as local domains in which the constitutive properties deviate significantly from their values in the undamaged material.

Another way of modeling a flaw such as a crack is through the imposition of appropriate boundary conditions at the edge of the flaw preventing the transmission of force across it. However, the use of boundary conditions will be problematic from a practical standpoint because it precludes the use of gradient-based techniques for calibration. Given the high dimensionality of the simulation parameter space incurred

by using implicit neural representations of the parameters of interest, gradient-based techniques are required to make the problem computationally tractable as there are no known gradient-free methodologies with competitive performance for high-dimensional optimization.

**Finite difference physics solver** We consider the use of a simulation-based calibration scheme built on a finite-difference (FD) implementation of a partial differential equation of the form

$$u_{tt} = f(u, u_x, u_y, u_{xx}, u_{yy}; \theta_f). \quad (4)$$

For the wave equation,

$$f = v^2(x, y; \theta_v)(u_{xx} + u_{yy}), \quad (5)$$

The required spatial derivatives are computed via a second-order central finite difference scheme. We use PyTorch (Paszke et al. 2019) to implement  $f$  as a convolution operator with fixed kernel matrix. Implementing the solver completely in PyTorch also results in the ability to easily incorporate neural representations of  $v$  and (Dirichlet) boundary conditions (Hoyer, Sohl-Dickstein, and Greydanus 2019; Sitzmann et al. 2020; Raissi, Perdikaris, and Karniadakis 2019) as well as to differentiate through the simulation via backward-mode automatic differentiation (backpropagation). While our implementation uses naïve automatic differentiation through the whole simulation similarly to (Degraeve et al. 2019), more memory-efficient differentiable solvers have been devised utilizing adjoint equations (de Avila Belbutte-Peres et al. 2018) which could be used to relieve the memory cost associated with scaling this method to finer-resolution simulations, though the time cost would not be significantly altered. Empirically, we find that the latter constitutes the tighter constraint for our problem.

**Physics-informed neural networks for solving inverse problems** The method of physics-informed neural networks is explained in detail in (Raissi, Perdikaris, and Karniadakis 2019). For physics-informed neural networks for inverse problems (henceforth referred to as iPINNs), one parameterizes the solution  $u(\cdot)$  to a differential equation with a neural network and supervises the network with the squared residual of the relevant partial differential equation as well as a set of observations on  $\hat{u}$ . This method uses gradient-based optimization to simultaneously optimize the parameters of the solution network as well as any other unknown parameters of the differential equation.

In the case of the current problem, (Shukla et al. 2020) define a solution net  $u(x, y, t)$  and a latent field net  $v(x, y)$  and supervise both on the heterogeneous wave equation of Eq. (3) as well as  $\mathcal{D}$  via a squared error loss on both. For completeness, we implement their method as a baseline for inversion when the physics are assumed to be known (i.e. Eq. (5) is used with only  $v(x, y)$  subject to learning).

**Bayesian Hidden Physics Models** Bayesian Hidden Physics Models (BHPMs) (Atkinson 2020) are compositional probabilistic models representing data from multiple experiments and their common underlying physics. The key differences between BHPMs and inverse PINNs are that the

physics are not assumed to be known *a priori* (up to some calibration parameters). The BHPM approach also extends the Deep Hidden Physics Models approach (Raissi 2018) by quantifying the epistemic uncertainty in the discovered partial differential operator, providing a quantitative answer to the informal question “What do the data tell us about the physics, and what *don’t* they tell us?”

Informally, BHPMs comprise “leaf” modules which are in charge of modeling information specific to a single experiment, while the root combines the functional representations of the leaves to learn a common physical governing law. Formally, let  $u : \mathbb{R}^{d_s+1} \rightarrow \mathbb{R}$  represent a scalar observable field in  $d_s$ -dimensional space and time as an implicit parameterization of the measurements in  $\mathcal{D}$ . Furthermore,  $u$  is assumed to obey some implicit relationship  $\Phi[u(\cdot)] = 0$ . In the context of the current problem, each leaf contains two functions:  $u(x, y, t)$  and  $a(x, y)$ , both of which are parameterized by multilayer perceptrons (MLPs) with sine activation functions. We also constrain  $a(\cdot) > 0$  by using an element-wise exponential activation function as its final layer. The former captures the measured deflection in space and time, and the latter is a latent field for which no observations are available. The two are related through the following differential equation:

$$u_{tt} = a^2(x, y)f(u, u_x, u_y, u_{xx}, u_{yy}), \quad (6)$$

Equation (6) may be thought of as a generalized wave equation; one recovers the familiar linear wave equation when  $a = v(x, y)$  and  $f = u_{xx} + u_{yy}$ , wherein  $a$  gains the physical significance of being the speed of sound in the medium. Imposing rotational invariance on  $f(\cdot)$  we obtain

$$u_{tt} = a^2(x, y)f(u, |\nabla u|, \Delta u), \quad (7)$$

We define the input to  $f$  as  $\psi = (u, |\nabla u|, \Delta u)$  for brevity.

Finally, reflecting our ignorance about the functional form of the physics, we place a Gaussian process prior over  $f(\cdot)$ :

$$f(\cdot) \sim \mathcal{GP}(\mu(\cdot), k(\cdot, \cdot)), \quad (8)$$

where  $\mu(\cdot)$  is a linear function and  $k$  is an exponentiated quadratic parametric kernel with automatic relevance determination (Rasmussen and Williams 2006).

Decomposing the right hand side of Eq. (7) is advantageous because it allows us to separate the properties of the particular microstructure of the specimen being studied from the general physical law; the former is represented by  $a$ , while the latter is captured in  $f$  and may be reused once learned. Furthermore, the multiplicative interaction between  $a^2$  and  $f$  is advantageous because it encourages an interpretable structure between the various components of the physics while the expressiveness of  $a$  and  $f$  allow us to discover rich interactions.

The structure of the right hand side of Eq. (7) implies that the product  $a^2 f$  is also *a priori* a Gaussian process-distributed:

$$a^2 f \sim \mathcal{GP}(a^2 \mu(\cdot), a^4 k(\cdot, \cdot)). \quad (9)$$

Due to the high volume of data passing through the  $f$  as well as the low complexity of the function it is ex-

pected to encode<sup>1</sup>, we approximate it using an uncollapsed sparse variational formulation (Hensman, Fusi, and Lawrence 2013), approximating the true posterior  $p(f(\cdot)|\mathcal{D}) \approx \int p(f(\cdot)|\mathbf{f}_u)q(\mathbf{f}_u)d\mathbf{f}_u$ , where the inducing variables  $\mathbf{f}_u \in \mathbb{R}^m$  associated with inducing locations  $\Psi_u \in \mathbb{R}^{m \times 3}$  are *a priori* distributed according to  $f(\cdot)$  and are given a Gaussian variational posterior.

As in (Atkinson 2020), we equip the measurement and physics targets with Gaussian likelihoods. Equation (2) is maximized with the variational parameters of  $f(\cdot)$ 's inducing points variables as  $\theta_f$  and the likelihood and MLP parameters as  $\theta_{\setminus f}$ .

## Related work

Bayesian calibration of physics codes to data has a long history, with the seminal work of Kennedy and O'Hagan (Kennedy and O'Hagan 2001). Other work (Bilionis and Zabaras 2013) showed how fully Bayesian inversion incorporating probabilistic surrogate models may be facilitated. More recently, the use of end-to-end differentiable physics engines for inverse problems in control has been demonstrated (de Avila Belbute-Peres et al. 2018; Degraeve et al. 2019; Rackauckas et al. 2020). Such methods enable the use of gradient-based inference and enable one to calibrate much higher-dimensional parameters, including implicit neural representations of unknown field quantities (Sitzmann et al. 2020). The Bayesian Hidden Physics Models approach (Atkinson 2020) combines the Bayesian aspects of earlier work with the powerful neural representations of the latter and can be used to conduct nonparametric inference on non-linear partial differential operators.

Regarding the application to microstructure characterization, Seidl (Seidl 2018) showed that the presence of cracks can be modeled as regions with locally-depressed wave speed. More recently, Shukla et al. (Shukla et al. 2020) showed that the problem of waveform inversion can be accomplished using the method of physics-informed neural networks (PINNs) (Raissi, Perdikaris, and Karniadakis 2019). The key difference between our work and these is that we do not specify the physics of the data *a priori* when using Bayesian Hidden Physics Models; instead, we allow ourselves to discover the physics from the data. Additionally, we capture the epistemic uncertainty in the physics due to the finite amount of data on which our inference is based, enabling us to identify what knowledge of the physics is still missing and enabling one to be able to carry out a principled adaptive experimental design algorithm.

The crack inversion problem has been tackled traditionally through physics-based finite element simulations. Fichtner has elaborated (Fichtner 2010) the aspects of full waveform tomography including methods for the numerical solution of the elastic wave equation, the adjoint method, the design of objective functionals and optimization schemes. Rao et al. (Rao, Ratassepp, and Fan 2016) proposed a guided wave

tomography method based on full waveform inversion for reconstruction of the remaining wall thickness in isotropic plates. The forward model is computed in the frequency domain by solving a full-wave equation in a two-dimensional acoustic model. Seidl et al. (Seidl 2018) elaborated the work to detect structural flaws and their position, dimension and orientation through acoustic full waveform inversion.

## Results

Code implementing the experiments can be found on GitHub.<sup>2</sup>

### Calibration of known physics to a pristine specimen

We first demonstrate the various modeling approaches on modeling the ‘‘pristine’’ specimen. Figure 1 shows a snapshot in time of the results calibrating the FD approach to data cropped to a  $1.5 \text{ mm} \times 2.5 \text{ mm} \times 1 \mu\text{s}$  window, while Fig. 2 shows the results for the iPINN approach applied to a  $3.75 \text{ mm} \times 2.5 \text{ mm} \times 3 \mu\text{s}$  window. For the FD approach we report the agreement between simulated and observed displacement fields as well as the inferred  $v(\cdot)$ . For the iPINN approach we additionally report on the residual of the physics equation since it is not enforced by construction. We see that both are able to reproduce the observed data and explain the data in terms of the wave equation physics that they assume.




Figure 1: Calibration of the FD approach to the dataset from the pristine specimen. Note that the left side of (b) is unreliable as the wave packet does not travel over it during the window used to calibrate the simulator.

<sup>1</sup>For example the wave equation is a linear equation in two scalar variables, and even differential equations whose solutions exhibit high complexity come from operators with comparatively simple structure.

<sup>2</sup><https://github.com/sdatkinson/BHPM-Ultrasound>

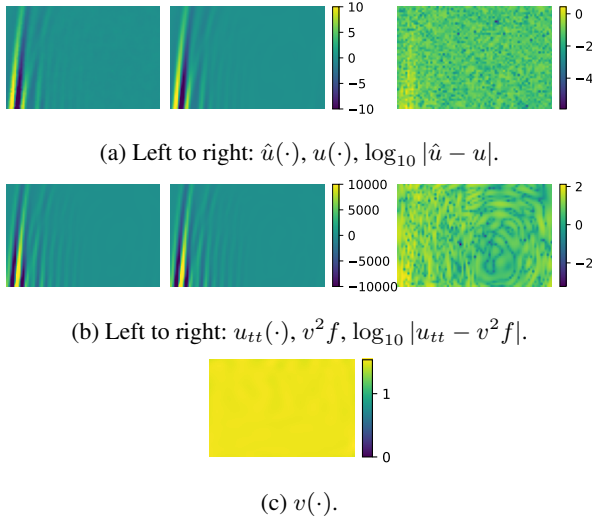


Figure 2: iPINN learned models for the observations and physics of the pristine specimen.

### Calibration of known physics to a cracked specimen

We now proceed to characterize a specimen with a localized flaw. We expect that this will show up as a corresponding region in which  $v$  is noticeably lower. This problem is challenging due to the details in the geometry of the flaw.

Figure 3 shows a snapshot in time of the calibration result using the FD method on the cracked specimen windowed to  $4.5 \text{ mm} \times 3 \text{ mm} \times 1.2 \mu\text{s}$ , and Fig. 4 shows compares the calibration using the iPINN model with data cropped only in time to  $5 \mu\text{s}$ . We were not able to calibrate as large an area for the FD solver due to its limitations related to discretization and the computational cost associated with the method.<sup>3</sup>

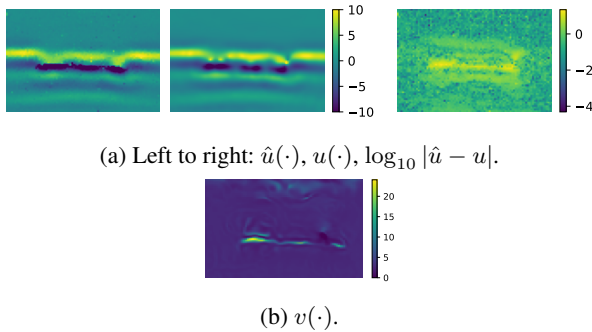


Figure 3: Calibration of the FD solver to the cracked specimen.

Comparing the FD and iPINN approaches, we see that there are some discrepancies in the neighborhood of the flaw. On one hand, the relationship between  $u$  and  $v$  is enforced to

<sup>3</sup>The FD approach parameters were selected to limit it to taking at most about an order of magnitude more time than the other approaches.

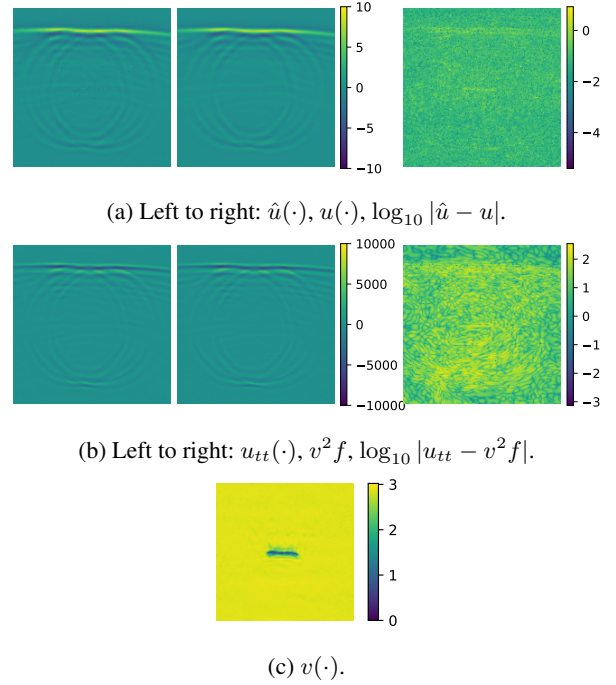


Figure 4: iPINN learned models for the observations and physics of the cracked specimen.

solver precision in the FD approach. By contrast, agreement in the physics is only enforced through a soft constraint in the case of the iPINN approach. Based on the observations that (1) the FD solver struggles to accurately reproduce the data in the vicinity of the crack and (2) that the physics are least in agreement in the vicinity of the crack for the iPINN approach, we infer that there may be missing physics in the standard wave equation that preclude its direct application to the case of flawed specimens. In the next sections, we explore this hypothesis by learning the physics from data using a Bayesian Hidden Physics Model with a larger physics hypothesis space rather than calibrating the postulated wave equation to data.

### Discovery of unknown physics from a pristine specimen

A Bayesian Hidden Physics model is constructed as discussed above and trained on data obtained from the pristine sample windowed to  $3.75 \text{ mm} \times 2.5 \text{ mm} \times 3 \mu\text{s}$ . Figure 5 shows the agreement obtained for the observation data and discovered physics. The root mean squared error for the observation function is  $1.35 \times 10^{-1}$  and the root mean squared residual of the learned physics using the posterior mean of  $f$  is  $1.44 \times 10^1$ .

### Characterization of flawed specimen using learned physics

We now transfer the learned physics from the pristine specimen to characterize the cracked specimen, windowed in time to  $5 \mu\text{s}$ . To do this, we initialize a second BHPM model, using the posterior on  $f(\cdot)$  learned from the first experiment. Only

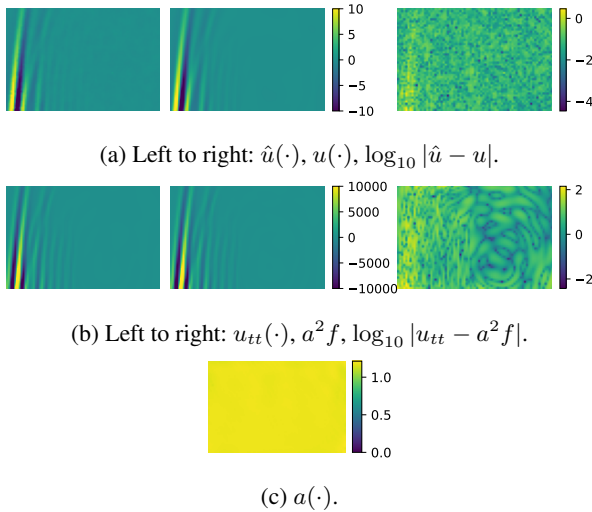


Figure 5: BHPM learned models for the observations and physics of the pristine specimen.

the MLPs for  $u$  and  $a$  are trained to characterize the sample. Figure 6 shows the agreement obtained for the observation data and discovered physics. The root mean squared error for the observation function is  $2.38 \times 10^{-1}$  and the root mean squared residual of the learned physics using the posterior mean of  $f$  is  $1.46 \times 10^2$ . Remarkably, we see that  $a(\cdot)$  takes on depressed values in the vicinity of the crack in the specimen, indicating that the backscattering in the observations can be traced back to the inhomogeneity in the material, while the far-field character is relatively homogeneous. Note also that the factors  $a$  and  $f$  comprising the right hand side are not individually uniquely identifiable since an arbitrary factor of  $\alpha$  could be multiplied and divided into  $a^2$  and  $f$ , respectively; therefore, the absolute magnitude of  $a$  in Figs. 5c and 6c is arbitrary and should not be directly compared with  $v(\cdot)$ . However, the relative magnitude within each field as well as the relative magnitude across experiments is meaningful due to the fact that  $f$  was held fixed.

## Conclusions

In this work, we considered the problem of characterizing the microstructure of metal specimens with and without flaws from data via physics discovery. We began by establishing agreement between two physics-based approaches using an end-to-end differentiable finite difference code as well as the inverse physics-informed neural networks approach. While both methods exhibited good agreement for pristine specimens, we found that they disagreed more in the case of a specimen with a crack flaw. Following this observation, we used Bayesian Hidden Physics Models to learn the physics from data from the pristine specimen, allowing for a larger hypothesis space than the standard wave equation. We then transferred the learned physics to the flawed specimen and characterized its microstructure through the lens of the learned physics. We found that the BHPM approach is able to identify the flaw in the specimen.

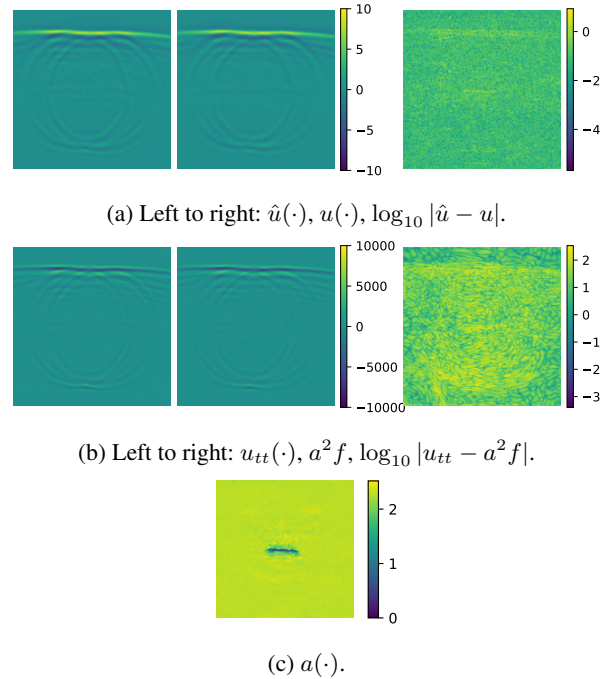


Figure 6: BHPM learned models for the observations and physics of the cracked specimen.

It is of general interest to develop an end-to-end scientific and engineering pipeline wherein one may begin by observing some novel behavior, collect measurements within a laboratory setting, deduce the relevant physical laws from those observations, then proceed with constructing a physically-accurate computational model of the discovered physics such that simulation may be leveraged for engineering design and analysis. All elements of this pipeline have been reduced to practice in the past. However, the most challenging element of the pipeline is that of discovery of novel physics from data; this is typically exclusively in the realm of human scientists. Our results using BHPM demonstrate that probabilistic machine learning techniques may expedite this process.

A number of questions of scientific interest persist, as do several opportunities for further engineering development based on our work. On the more fundamental side, the probabilistic side of physics discovery through machine learning is in its infancy. However, the recent advance of the BHPM approach stands to unite the large body of Bayesian inquiry to more open-ended exploration and discovery in physics; in particular, questions of optimal experimental design may now be treated rigorously, as the BHPM approach places questions of physics knowledge with an information-theoretic context. On the side of engineering development and capability, it is of interest to incorporate the learned physics into novel computational models that can be brought to bear to explore the space of material behaviors under various heterogeneities and external actions. The construction of physics-based codes and surrogates therein will lead the way to efficient sensing methodologies capable of real-time characterization and monitoring of materials and structures in engineering systems.

## Acknowledgments

The authors thank the United States Air Force Research Laboratory for providing the data analyzed in this work for public release under Distribution A as defined by the United States Department of Defense Instruction 5230.24. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR00111990032.

We thank the developers of the software from which our research code development has benefitted, including H5Py (Collette 2013), JAX (Bradbury et al. 2018), Matplotlib (Hunter 2007), NumPy (Oliphant 2006; Van Der Walt, Colbert, and Varoquaux 2011), Pandas (McKinney et al. 2011), PyTorch (Paszke et al. 2019), Scikit-learn (Pedregosa et al. 2011), SciPy (Virtanen et al. 2020), and tqdm (da Costa-Luis 2019).

## References

- Atkinson, S. 2020. Bayesian Hidden Physics Models: Uncertainty Quantification for Discovery of Nonlinear Partial Differential Operators from Data. *arXiv preprint arXiv:2006.04228*.
- Bilonis, I.; and Zabarar, N. 2013. Solution of inverse problems with limited forward solver evaluations: a Bayesian perspective. *Inverse Problems* 30(1): 015004.
- Blei, D. M.; Kucukelbir, A.; and McAuliffe, J. D. 2017. Variational inference: A review for statisticians. *Journal of the American statistical Association* 112(518): 859–877.
- Bradbury, J.; Frostig, R.; Hawkins, P.; Johnson, M. J.; Leary, C.; Maclaurin, D.; and Wanderman-Milne, S. 2018. JAX: composable transformations of Python+NumPy programs. URL <http://github.com/google/jax>.
- Collette, A. 2013. *Python and HDF5*. O'Reilly.
- da Costa-Luis, C. O. 2019. tqdm: A fast, extensible progress meter for python and cli. *Journal of Open Source Software* 4(37): 1277.
- de Avila Belbute-Peres, F.; Smith, K.; Allen, K.; Tenenbaum, J.; and Kolter, J. Z. 2018. End-to-end differentiable physics for learning and control. In *Advances in Neural Information Processing Systems*, 7178–7189.
- Degrave, J.; Hermans, M.; Dambre, J.; et al. 2019. A differentiable physics engine for deep learning in robotics. *Frontiers in neurorobotics* 13: 6.
- Fichtner, A. 2010. *Full seismic waveform modelling and inversion*. Springer Science & Business Media.
- Hensman, J.; Fusi, N.; and Lawrence, N. D. 2013. Gaussian processes for big data. *arXiv preprint arXiv:1309.6835*.
- Hoyer, S.; Sohl-Dickstein, J.; and Greydanus, S. 2019. Neural reparameterization improves structural optimization. *arXiv preprint arXiv:1909.04240*.
- Hunter, J. D. 2007. Matplotlib: A 2D graphics environment. *Computing in Science & Engineering* 9(3): 90–95. doi:10.1109/MCSE.2007.55.
- Kennedy, M. C.; and O'Hagan, A. 2001. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63(3): 425–464.
- McKinney, W.; et al. 2011. pandas: a foundational Python library for data analysis and statistics. *Python for High Performance and Scientific Computing* 14(9): 1–9.
- Oliphant, T. E. 2006. *A guide to NumPy*, volume 1. Trelgol Publishing USA.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, 8026–8037.
- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12: 2825–2830.
- Rackauckas, C.; Ma, Y.; Martensen, J.; Warner, C.; Zubov, K.; Supekar, R.; Skinner, D.; and Ramadhan, A. 2020. Universal differential equations for scientific machine learning. *arXiv preprint arXiv:2001.04385*.
- Raissi, M. 2018. Deep hidden physics models: Deep learning of nonlinear partial differential equations. *The Journal of Machine Learning Research* 19(1): 932–955.
- Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* 378: 686–707.
- Rao, J.; Ratssepp, M.; and Fan, Z. 2016. Guided wave tomography based on full waveform inversion. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* 63(5): 737–745.
- Rasmussen, C. E.; and Williams, C. K. 2006. *Gaussian processes for machine learning*. MIT press Cambridge, MA.
- Seidl, R. 2018. *Full Waveform Inversion for Ultrasonic Nondestructive Testing*. Ph.D. thesis, Technische Universität München.
- Shukla, K.; Di Leoni, P. C.; Blackshire, J.; Sparkman, D.; and Karniadakis, G. E. 2020. Physics-informed neural network for ultrasound nondestructive quantification of surface breaking cracks. *arXiv preprint arXiv:2005.03596*.
- Sitzmann, V.; Martel, J. N.; Bergman, A. W.; Lindell, D. B.; and Wetzstein, G. 2020. Implicit Neural Representations with Periodic Activation Functions. *arXiv preprint arXiv:2006.09661*.
- Van Der Walt, S.; Colbert, S. C.; and Varoquaux, G. 2011. The NumPy array: a structure for efficient numerical computation. *Computing in Science & Engineering* 13(2): 22.
- Virtanen, P.; Gommers, R.; Oliphant, T. E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P;

Weckesser, W.; Bright, J.; van der Walt, S. J.; Brett, M.; Wilson, J.; Jarrod Millman, K.; Mayorov, N.; Nelson, A. R. J.; Jones, E.; Kern, R.; Larson, E.; Carey, C.; Polat, İ.; Feng, Y.; Moore, E. W.; Vand erPlas, J.; Laxalde, D.; Perktold, J.; Cimrman, R.; Henriksen, I.; Quintero, E. A.; Harris, C. R.; Archibald, A. M.; Ribeiro, A. H.; Pedregosa, F.; van Mulbregt, P.; and Contributors, S. . . 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17: 261–272. doi:<https://doi.org/10.1038/s41592-019-0686-2>.