

# Physics-Guided Data-Driven Seismic Inversion:

## *Recent Progress and Future Opportunities in Full Waveform Inversion*

Youzuo Lin, James Theiler, and Brendt Wohlberg

### Abstract

The goal of *seismic inversion* is to obtain subsurface properties from surface measurements. Seismic images have proven valuable, even crucial, for a variety of applications, including subsurface energy exploration, earthquake early warning, carbon capture and sequestration, estimating pathways of subsurface contaminant transport, *etc.* These subsurface properties (such as wave speed, density, and elastic velocities) influence the transmission of seismic waves through the subsurface media, and well-understood physics models (so-called “forward models”) can be used to predict what surface measurements would be made, for any given subsurface configuration. Seismic inversion is the inverse problem: given actual surface measurements, infer what subsurface configuration would give rise to those measurements. Like most inverse problems, seismic imaging is ill-posed, meaning many different subsurface configurations can give rise to the same surface measurements. Iterative optimization algorithms for the inverse problem are typically very computationally expensive because they require many evaluations of the forward model, which is itself computationally expensive. A further challenge is the different sensitivity of subsurface properties to the seismic data; density for example is more difficult to accurately infer than is P-wave velocity.

But recent advances in algorithms and computing provide an opportunity for remarkable progress in seismic inversion, and efficient solutions to previously infeasible problems have been obtained using data-driven approaches (such as the deep learning methods that were developed primarily for problems in computer vision). The excellent performance of learning-based methods arises from its ability to exploit large amounts of high-quality training data, without the need for hand-designed features. Unlike computer vision, however, seismic inversion is not a data-rich domain. There is a relatively small amount of field data in existence due to the high cost of acquisition, and as a result of its commercial value, a very limited amount is publicly available. To alleviate the data scarcity issue and improve model generalization, there has been growing interest in combining physics knowledge with machine learning for solving seismic inversion problems.

Y. Lin, J. Theiler, and B. Wohlberg are with Los Alamos National Laboratory, Los Alamos, NM, 87544 USA

Manuscript was received March, 2022.

This review will survey methods for incorporating physics knowledge with machine learning (primarily deep neural networks) to solve computational seismic inversion problems. We will provide a structured framework of the existing research in the seismic inversion community, and will identify technical challenges, insights, and trends.

### Index Terms

Scientific Machine Learning, Physics Guided Neural Network, Seismic Inversion, Subsurface Characterization, Computational Imaging, Inverse Problems

## I. INTRODUCTION

WHILE we can't directly observe the geology of the earth beneath our feet, knowledge of its properties is vital for a variety of applications, including knowing where to build or drill for water or hydrocarbons. From civil infrastructure to energy exploration, characterizing subsurface geology is crucial, but since exhaustive sampling of the earth's interior is nearly impossible, inference of unknown subsurface properties primarily relies on indirect and limited geophysical measurements. Seismic inversion attempts to reconstruct an image of the subsurface from measurements of natural or artificially produced seismic waves that have traveled through the subsurface medium. We emphasize surface seismic measurements in the examples here, but a vertical seismic profile (VSP) can also be obtained from sensors within a borehole. This provides complementary information – less spatial coverage but potentially higher spatial resolution – compared to the surface seismic approach. A forward model describes how the observations depend on the subsurface properties, while the inverse problem is the inference of those properties from the observations. The forward model of seismic wave propagation is governed by the wave equation, in the form of a partial differential equation (PDE). A linear approximation of the inversion scheme with high-frequency approximation of ray tracing can be utilized to simplify the problem, at the cost of loss of accuracy and resolution, leading to travel-time inversion methods [1]. Full-waveform inversion (FWI), which considers full waveform information including both kinematic (amplitude) and dynamic (phase) information, addresses the full non-linear problem, providing superior inversion accuracy and resolution [2], [3], but at a considerably greater computational cost. Given the large volume of publications in this field, our review focuses on non-linear seismic FWI techniques over other seismic inverse problems (such as travel-time tomography) based on the following considerations:

- **Superior Performance.** FWI leverages extensively more waveform physics than other seismic inversion techniques, resulting in significantly improved imaging results. The technique has been




Fig. 1: (a) Physics-guided data-driven seismic imaging. The wave equation is the governing physics. (b) Trends of physics-guided machine learning techniques in computational seismic imaging. Six major technical categories can be identified out of a total 107 papers from 2017 to 2021 collected from the Web of Science and Google Scholar.

widely applied to various subsurface applications ranging from reservoir-scale subsurface velocity building [4] to global-scale tomography of Earth’s crust and mantle [5].

- **Significant Technical Impact.** FWI faces similar technical challenges to many other ill-posed inverse problems. Our focus here is on how to leverage the underlying governing wave physics (as shown in Fig. 1a), but the techniques discussed here can be not only extendable to other wave imaging problems but also provide potential solutions to a broader context of PDE-governed inverse problems.
- **Active Research.** FWI is one of the most active research topics in subsurface geophysics. A variety of new computational methods have emerged recently (see Fig. 1b) for solving the FWI problem. This forms an excellent foundation for us to provide a broad and in-depth review of the existing literature.

Seismic full-waveform inversion is a nonlinear inverse problem. Approaches for solving this problem can be categorized into two groups: physics-based methods and machine learning-based methods. Physics-based methods primarily rely on nonlinear optimization techniques that are computationally intensive and highly sensitive to data coverage<sup>1,2</sup>. Their strong dependence on the initial estimate, and the difficulty

<sup>1</sup>Data coverage: the spatial extent (and receiver/source density) to which a subsurface region is covered by a seismic survey.

<sup>2</sup>We provide a summary of the seismic terminologies used in the review (see Table II).

of choosing the best problem parameterization<sup>3</sup> also make those methods challenging in applications. Furthermore, some additional computational issues that arise from real applications have also attracted attention; these include efficient 3D inversion, uncertainty assessment, inelastic domain, and anisotropy [3]. Once developed, however, they can be robust and generalize quite naturally to new datasets. In part to alleviate the computational expense of physics-driven methods, data-driven inversion techniques have recently been proposed [6], [7], [8], [9] and reviewed [4]. The data-driven techniques can be extremely efficient once fully trained, but the generalizability of the data-driven approaches is limited by the size and range of the training set. In addition, there is a relatively small amount of field data in existence due to the high cost of acquisition, and as a result of its commercial value, a very limited amount is publicly available. Furthermore, because data-driven models learn from data directly without leveraging domain knowledge, those models are prone to unreasonable or unrealistic predictions that do not conform to the physical mechanism. Efforts have been made to combine the physical consistency and generalization ability of physics-based modeling with the computational efficiency of data-driven approaches; these hybrid algorithms are collectively referred to as physics-guided machine learning (PGML). There is ongoing work in the seismic inversion community to develop new models that combine physics-driven methods with machine learning techniques. This review paper will focus on this new trend. The strengths and limitations of these techniques will be discussed based on some numerical examples.

Several review articles relevant to this topic have recently been published. We summarize a number of them in Table I, and categorize them into three groups: physics-based seismic imaging techniques, machine learning methods (ML), and physics-informed machine learning (PIML).

- **Physics-based.** Virieux and Operto [3] provide a thorough overview of physics-based FWI techniques in exploration geophysics, while Tromp [5] gives a different perspective in an overview focusing on applications of FWI for global seismology problems. In addition to seismology, FWI also finds application to medical imaging, including in brain imaging [17] and breast imaging [18]. A good review of those topics can be found in a recent thesis [10].
- **Machine Learning.** We have found four particularly noteworthy review papers that focus on data-driven methods. Bergen et al. [11] provide a summary of machine learning models for solving general solid Earth geosciences problems ranging from earthquake physics to data visualization, with a focus on existing supervised and unsupervised ML models developed for subsurface applications. Gregory

<sup>3</sup>There are a number of different possible ways of representing properties of the subsurface, including wave propagation velocity, slowness (the reciprocal of velocity) and squared slowness, the choice of which influences the objective functional gradient and Hessian, which typically play a critical role in the optimization.

Reference	Year	Application Domain	Physics	ML	PGML
Virieux and Operto [3]	2009	Exploration Geophysics	✓		
Agudo [10]	2018	Medical Imaging	✓		
Tromp [5]	2020	Global Geophysics	✓		
Bergen et al. [11]	2019	General Applications in Solid Earth		✓	
Gregory et al. [12]	2020	Computational Imaging and Inverse Problems		✓	
Yu and Ma [13]	2021	General Applications in Solid Earth		✓	
Adler et al. [4]	2021	Computational Seismic Inverse Problems		✓	
Willard et al. [14]	2020	Engineering and Environmental Problems			✓
Karniadakis et al. [15]	2021	Scientific and Engineering Problems			✓
Cuomo et al. [16]	2022	Scientific and Engineering Problems			✓

TABLE I: A summary of relevant review articles in the topics of physics-based seismic FWI for exploration geophysics [3], medical imaging [10], and global geophysics [5]; pure ML approaches for general geoscience [11], [13], computational imaging [12], and seismic inverse problems [4]; and PGML for environmental problems [14], and for broad scientific and engineering inverse problems [15], [16].

et al. [12] provide an overview on the broad topics of deep learning for computational imaging and inverse problems, concentrating mainly on methodology development rather than on specific imaging problems. In contrast, Yu and Ma [13] explore deep neural networks, and emphasize a broad range of applications in geophysics, including geophysical data processing, interpretation, earthquake detection, and tomography. Adler et al. [4] provide an overview of recent machine learning-based methods for solving different computational seismic inverse problems including velocity, impedance, reflectivity model building, and seismic bandwidth extension. The focus of their paper is mostly on pure data-driven deep learning approaches with applications to exploration geophysics.

- **PGML.** Physics-guided machine learning (PGML) is a relatively new technique that has, thus far, seen quite limited application in geoscience. Among existing PGML models, the physics-informed neural network (PINN) [19] is one of the most well-known approaches. There are some reviews on PGML for general topics [20], [15], but, none of them focuses specifically on computational seismic imaging. Willard et al. [20] discuss the use of PGML in engineering and environmental problems. Karniadakis et al. [15] provide a more high-level overview of PINN with applications to inverse problems in different domains such as biophysics, quantum chemistry, material sciences, and molecular simulations. A very recent review paper [16] provides a comprehensive review of the literature on PINN and its variants, discussing their related advantages, disadvantage, and existing

Terminology	Description in Plain Language
Seismic Data Coverage	The spatial extent and receiver/source density to which a subsurface region is covered by a seismic survey.
Velocity Maps	Spatial maps of the local speed of the seismic waves through the subsurface medium. There may be multiple maps, since different kinds of waves can have different speeds.
Well Logs	Well logs present a record of one or more physical measurements (such as acoustic velocity, electrical resistivity, gamma ray, etc.) as a function of depth in a borehole.
Reverse Time Migration	Reverse time migration (RTM) is a seismic imaging technique focusing on seismic impedance contrasts, that is able to estimate subsurface structural information, such as the locations of layer boundaries. One of the impacting factors of RTM imaging accuracy is the quality of input velocity maps.
Prestack Data	Prestacking is a seismic data pre-processing technique and prestack data is the data before stack/migration.
Post-stack Data	Post-stack data refers to the data that is processed after stacking and results in a subsurface structural image, which is in the time domain. Here, the stacking operation is the process to sum seismic traces to improve the signal-to-noise ratio.

TABLE II: A summary of seismic terminology used in this review.

technical challenges.

The remainder of this paper is organized as follows. Section II (Current Research and Development) provides a high-level overview of the physics-based and data-driven computational seismic imaging methods. Section III (Physics-Informed Data-Driven Computational Seismic Imaging) gives a summary of existing hybrid models incorporating physics with machine learning techniques. Section IV (Challenges and Future Research Opportunities) provides a perspective of the existing challenges and some future opportunities. Finally, Section V (Conclusions) summarizes our paper.

## II. CURRENT RESEARCH AND DEVELOPMENT

### A. Physics-Based Seismic Imaging Techniques

Physics-driven methods infer subsurface properties by directly employing the governing physics and equations. Seismic waves are mechanical perturbations that travel in the medium at a speed governed by the acoustic/elastic impedance of the material through which they are traveling. In the time domain, the

acoustic wave equation is given by

$$\nabla^2 p(\mathbf{r}, t) - \frac{1}{c^2(\mathbf{r})} \frac{\partial^2}{\partial t^2} p(\mathbf{r}, t) = s(\mathbf{r}, t), \quad (1)$$

where  $t$  denotes time,  $\mathbf{r} = (x, z)$  represents the spatial location in Cartesian coordinates ( $x$  is the horizontal direction and  $z$  is the depth),  $\nabla^2$  is the Laplacian operator, and  $c(\mathbf{r})$  is the velocity map (that is, the speed of the wave through the medium at position  $\mathbf{r}$ ). Here,  $p(\mathbf{r}, t)$  is the pressure wavefield, and  $s(\mathbf{r}, t)$  is the source term that specifies the location and time history of the source. In the context of this forward model, if we know the velocity map  $c(\mathbf{r})$  and the source term  $s(\mathbf{r}, t)$ , we can use Eq. (1) to solve for the pressure field  $p(\mathbf{r}, t)$ . Thus, we can write

$$\mathbf{d} = f(\mathbf{m}), \quad (2)$$

to express the deterministic relationship between the model  $\mathbf{m}$ , which encapsulates the subsurface velocity map  $c(\mathbf{r})$ , and the observed data  $\mathbf{d}$ , which typically consists of time histories of the pressure wavefield  $p(\mathbf{r}, t)$  at a discrete set of locations  $\mathbf{r}$ , usually on the surface. The goal of the inverse problem is to infer the subsurface model  $\mathbf{m}$  given observed data  $\mathbf{d}$ . This can be posed as a minimization problem [21]

$$E(\mathbf{m}) = \min_{\mathbf{m}} \left\{ \|\mathbf{d} - f(\mathbf{m})\|_2^2 + \lambda \mathcal{R}(\mathbf{m}) \right\}, \quad (3)$$

where  $\|\mathbf{d} - f(\mathbf{m})\|_2^2$  is the data fidelity term with  $\|\cdot\|_2$  denoting the  $\ell_2$  norm, and  $\mathcal{R}(\cdot)$  is the regularization term with regularization parameter  $\lambda$ . Two of the most widely used regularization terms are Tikhonov regularization, which corresponds to setting  $\mathcal{R}(\mathbf{m}) = \|H\mathbf{m}\|_2^2$ , where  $H$  is a high-pass filtering operator, and total-variation (TV) regularization, which corresponds to setting  $\mathcal{R}(\mathbf{m}) = \|\mathbf{m}\|_{\text{TV}}$ , where

$$\|\mathbf{m}\|_{\text{TV}} = \sum_{1 \leq i, j \leq n} \sqrt{|(\nabla_x \mathbf{m})_{i,j}|^2 + |(\nabla_z \mathbf{m})_{i,j}|^2}, \quad (4)$$

with  $(\nabla_x \mathbf{m})_{i,j} = \mathbf{m}_{i+1,j} - \mathbf{m}_{i,j}$  and  $(\nabla_z \mathbf{m})_{i,j} = \mathbf{m}_{i,j+1} - \mathbf{m}_{i,j}$ . Tikhonov regularization [1] is simpler to implement computationally, but tends to blur sharp interfaces between regions, which are better preserved by TV regularization [21].

Solving the minimization in Eq. (3) can be numerically challenging or even infeasible due to the ill-posed nature of the inversion problem (i.e. even small noise levels in the measurements  $\mathbf{d}$  can yield large uncertainties in the model  $\mathbf{m}$ ). For the FWI problem, this can arise from limited data coverage (too few measurements of the pressure wavefield). Another notorious issue in solving FWI is *cycle skipping*<sup>4</sup> [3], which causes erroneous model updates, leading to an incorrect local minimum. Many approaches have been proposed to alleviate the aforementioned computational issues. These includes

<sup>4</sup>Cycle-skipping occurs when the phase match between the two data sets is greater than half a wavelength.




Fig. 2: An illustration of InversionNet, a data-driven seismic waveform inversion solver [7].

advanced regularization [21], dynamic warping techniques [22], incorporation of prior information [23], multiscale inversion approaches [24], and preconditioning techniques [25]. Another issue with current computational methods for solving FWI is the high computational cost. Most of the current techniques are based on first-order or second-order optimization techniques. Both of these types of methods are expensive since they require computation of the gradient of the function Eq. (3), which requires a forward wave propagation, a backward wave propagation, and full-wavefield cross-correlation. A line search strategy, requiring additional steps of forward wave propagation, is also necessary. To reduce the computational cost, dimensionality reduction techniques can be utilized to reduce the data volume by replacing many sequential sources with much fewer phase-encoded simultaneous sources [26].

### B. Data-Driven Seismic Imaging Techniques

We begin this section with the remark that *purely* data-driven techniques are unlikely ever to work for seismic imaging; there just isn't enough data. A lot of wells have been drilled over the last century, and so we have a reasonable idea about what the subsurface is like in general. But for a given site, there is precious little in the way of direct observations of the subsurface. When we speak of data-driven methods being applied to the seismic imaging problem, we are virtually always speaking of simulated data. As with traditional “physics-based” methods described in the previous subsection, we assume we know the governing equation for acoustic waves, and that we can readily simulate the function  $f(\cdot)$  described in Eq. (2). Given this, the key idea for data-driven methods is to create many samples of



image pairs: of velocity model  $\mathbf{m}$ -images along with matched samples of waveform data  $\mathbf{d}$ -images. It is vital that the  $\mathbf{m}$ -images are representative of the subsurface. Geostatistical methods [27] can be useful in creating ensembles of these  $\mathbf{m}$ -images that are geologically meaningful. But once we have a collection of  $\mathbf{m}$ -images, we can use  $\mathbf{d} = f(\mathbf{m})$  to create the associated  $\mathbf{d}$ -images. With all of this simulated data, we can at last apply some data-driven analysis. A large image-to-image deep neural net (as shown in Fig. 2) is used to learn a direct mapping of  $\mathbf{d}$  to  $\mathbf{m}$ , which can finally be used to map an actual data observation  $\mathbf{d}_o$  to an estimate of the subsurface map. Mathematically, this can be posed as

$$\begin{aligned} \widehat{\mathbf{m}}(\mathbf{d}) &= g_{\theta^*}(\mathbf{d}) \\ \text{s.t. } \theta^*(\Phi_s) &= \arg \min_{\theta} \sum_{(\mathbf{m}_i, \mathbf{d}_i) \in \Phi_s} \mathcal{L}(g_{\theta}(\mathbf{d}_i), \mathbf{m}_i), \end{aligned} \quad (5)$$

where:  $\theta$  represents the trainable weights in the inversion network  $g_{\theta}(\cdot)$ ;  $\mathcal{L}(\cdot, \cdot)$  is a loss function; and  $\Phi_s$  is the training dataset with paired samples of  $(\mathbf{m}_i, \mathbf{d}_i)$ .

A distinctive feature of this approach is that, in contrast with traditional inversion, it produces a “whole function”  $g_{\theta^*}(\cdot) \approx f^{-1}(\cdot)$ , not just a single inversion  $\widehat{\mathbf{m}}_{\text{obs}} \approx f^{-1}(\mathbf{d}_{\text{obs}})$ . Nonetheless, experience with data-driven methods indicates that they can be trained in substantially less time than it takes for a single traditional inversion [7]; and once trained, the application to a specific observation  $\mathbf{d}_{\text{obs}}$  is nearly instantaneous. The reason for this efficiency is not fully understood, but some factors that potentially contribute include: the simulations are only of the forward model, and no gradients (or Hessians) are computed; the simulations are restricted to plausible models, so no time is spent (as may be the case for traditional optimizers) in implausible regions of solution space; traditional optimization is inherently serial, whereas the many simulations can be readily parallelized, and neural network architectures naturally lend themselves to efficient massive parallelization. Further, because the function  $f^{-1}(\mathbf{d})$  is learned directly and implicit regularization is imposed via training on a large volume of simulations, we do not see the ill-posedness that traditional inversion often exhibits; on the other hand, solutions tend to look like the models  $\mathbf{m}$  that they were trained on, so the generalization ability of this approach is limited.

Most recent explorations have focused on data-driven end-to-end reconstruction. For instance, GeoDNN [8] used a fully connected network with 8 layers for FWI. InversionNet [7] introduced an encoder-decoder network with modern convolutional neural network (CNN) features. ModifiedFCN [9] employed a fully connected convolutional network (FCN) design and SeisInvNet [6] enhanced each seismic waveform with auxiliary knowledge from neighborhood traces for better spatial correspondence.

### III. PHYSICS-GUIDED DATA-DRIVEN SEISMIC INVERSION

In this section we discuss a variety of computational techniques that have been developed for incorporating physics in seismic imaging methods.

#### A. Physics Knowledge

Physics knowledge can be represented in many different ways. What makes seismic imaging problems fundamentally different from those of the computer vision community are the physical laws underlying those problems. These physical laws are “invariant” in the sense that the spatial and temporal dynamics can be characterized by those laws throughout the physical process. Here we provide two different representations of physics knowledge that has been utilized to enhance the machine learning model.

- 1) **Wave Physics.** The most directly relevant physics is expressed by the governing equations underlying the imaging process. For seismic imaging, this is the partial differential equation in Eq. (1) that describes the propagation of acoustic waves through the subsurface.
- 2) **Physical Properties of the Solution Space.** To the extent that we understand the geophysics of the subsurface, we can constrain the solution space to be consistent with those physical properties, and we can more effectively choose “plausible” samples from the solution space to train our neural networks. Among these properties are the tendencies for materials to stratify into distinct layers, the existence of faults as distinct geological formations, common tectonic regimes that create subsurface formations, and the overall variability of densities and wave speeds in subsurface material. We can also make use of dynamical properties of the subsurface phenomena that we are trying to measure. Consider CO<sub>2</sub> sequestration as one example. After injection of super-critical CO<sub>2</sub> into the subsurface, the CO<sub>2</sub> plume will migrate over time, meaning that the spatial spreading of CO<sub>2</sub> should gradually increase over time. If a model can be designed so that the images are constrained to exhibit dynamic spreading (and not, say, compaction) of the CO<sub>2</sub> plume, then this would improve the overall performance [28].

We provide a taxonomy of physics-informed learning-based seismic inversion techniques (see Fig. 3) with two major categories of approaches: one being built on ML models (blue hexagons) and the other built on physics models (green hexagons). There are three groups of techniques within each category. We will discuss all six groups in the following sections.

#### B. Machine Learning Models Incorporating Physics

In this section, we describe methods identified with blue hexagons in Fig. 3. These models are all fundamentally data-driven (in most cases, they are deep neural networks), but in each case with some

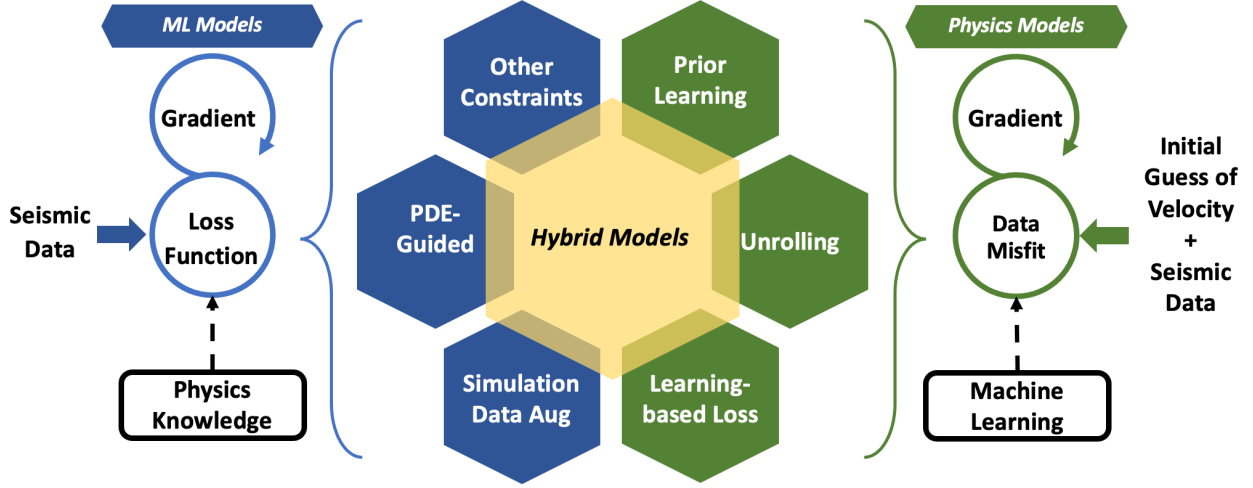


Fig. 3: An illustration of PGML-based seismic imaging techniques, which can be grouped into two categories: those built on ML models (blue hexagons) and those built on physics models (green hexagons).

aspect of the appropriate physics incorporated. In particular, we will discuss how to embed physics knowledge through data augmentation techniques, how to constrain the learning model using the governing wave equation, and how to learn underlying solution space distribution to help with the inversion.

1) *Simulations and Data Augmentation:* As noted previously in Section II-B, we can leverage our knowledge of the wave physics to simulate large quantities of paired data  $(m, d)$  where  $m$  is a plausible subsurface velocity map and  $d$  is the data that would be observed if  $m$  accurately modeled the subsurface. What we did not say in Section II-B, however, is what we meant by “plausible”, nor how a large sample of plausible subsurface maps might be generated. In general, doing this requires domain knowledge of the subsurface rock physics, and the creation of a large idealized set of subsurface structure maps (to be used as training sets) is commonly known as “velocity model generation” in the subsurface geophysics community.

While employing simulations generated from wave physics to incorporate knowledge can be effective, it may not ultimately be scalable, because it often leads to very large training sets – for example, more than 60,000 velocity maps were created for training in [7] – that still may not encompass the full variability of what might be in the subsurface. One approach is to use auxiliary knowledge of the subsurface properties to reduce this variability. Well logs<sup>5</sup> from drilled wells has been used to build a large-scale subsurface

<sup>5</sup>Well logs present a record of one or more physical measurements (such as acoustic velocity, electrical resistivity, gamma ray intensity, etc.) as a function of depth in a borehole.

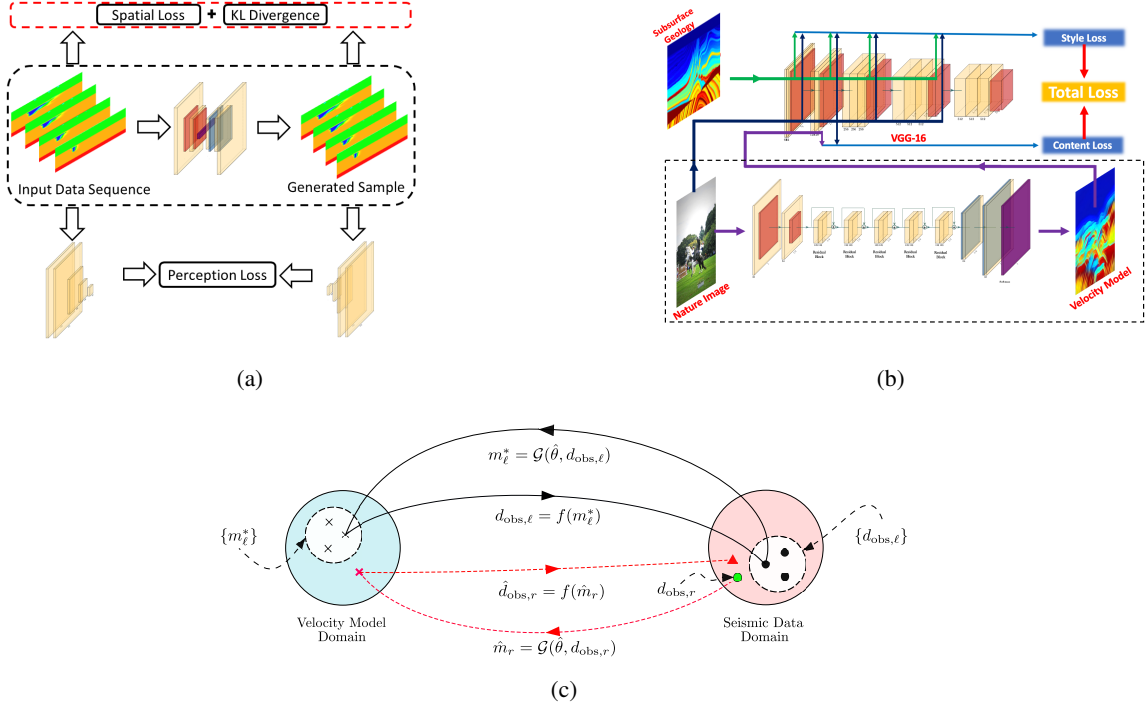


Fig. 4: Schematic illustration of three physics-informed data augmentation techniques to generate physically realistic samples. (a) VAE model with physics-informed loss to synthesize velocity maps [31]. (b) Style transfer model to leverage existing natural images to generate velocity maps [32]. (c) An adaptive active learning method to augment training set with unlabeled seismic data [33].

data set for training a deep learning model [29], and pseudo-wells based on the well statistics (such as porosity, saturation, mineralogy, and thickness) have also been created to train the neural network models [30]. Additionally, generative models (such as the variational auto-encoder and the generative adversarial network) have been leveraged to synthesize more realistic simulations. A variational auto-encoder with physics-informed regularization (shown in Fig. 4a) was designed to synthesize realistic velocity maps [31]. Domain adaptation is another technique that was shown to be effective to overcome the data scarcity issue. A style transfer network (see (Fig. 4b) was developed to convert from natural images to physically meaningful subsurface velocity maps [32].

It is possible, albeit less straightforward, to augment the training set from the seismic data perspective rather than just produce more simulations from a broader variety of velocity maps. A couple of authors have shown the benefits of enhancing supervised training with this kind of data augmentation [33], [34]. Specifically, Gómez et al. [33] developed adaptive data augmentation techniques based on active learning, where physically realistic simulated seismic data are created to enrich the existing training




Fig. 5: (a) Ground truth. Inversions and errors obtained using (b) physics-based FWI (0.069, 0.972), data-driven model trained on (c) the *large* and *medium* subsets (0.0620, 0.990), (d) *large* and *augmented medium* subsets (augmented once) (0.0134, 0.993), and (e) *large* and *augmented medium* subsets (augmented twice) (0.0122, 0.994). Error is in the format of (MAE, SSIM).

set. Ideas of domain adaptation have also been applied to enrich the quality of seismic data in training set. Alkhalifah et al. [34] construct a high-quality synthetic seismic dataset with more realistic features from the target seismic data (such as noise and artifacts). Their example applications on passive seismic imaging demonstrate the efficacy of their proposed techniques.

As a demonstration, we provide a brief overview of the technique developed in [33]. Four steps (illustrated in Fig. 4c) are proposed including:

- 1) Estimate an approximate solver  $\mathcal{G}(\hat{\theta}, \mathbf{d}_{\text{obs}})$ ;
- 2) Generate approximate velocity maps from unlabeled data  $\hat{\mathbf{m}}_{\text{obs}} = \mathcal{G}(\hat{\theta}, \mathbf{d}_{\text{obs}})$ ;
- 3) Create seismic data using forward model  $\hat{\mathbf{d}}_{\text{obs}} = f(\hat{\mathbf{m}}_{\text{obs}})$ ;
- 4) Add new pairs to the original training set.

The augmented dataset plays a key role in model accuracy because it will carry useful wave physics information in the regime of data-space that is closer to the  $\mathbf{d}_{\text{obs}}$  of interest. The full augmentation process can be applied in an iterative fashion by re-training the approximate solver  $\mathcal{G}(\hat{\theta}, \mathbf{d}_{\text{obs}})$  based on the extended training set in order to generate new approximate velocity maps  $\hat{\mathbf{m}}_{\text{obs}}$ . This approach allows further refinement of the mapping between velocity and seismic subdomains. Iterating from Steps 1 to 4 will produce a consistent refinement of the solver  $\mathcal{G}(\hat{\theta}, \mathbf{d}_{\text{obs}})$ . In Fig. 5, we further provide visualization of the inverted velocity maps using both physics-based FWI and data-driven approaches as proposed in [33] to characterize tiny CO<sub>2</sub> leakage plumes at the Kimberlina sequestration site [35]. We observe that the network trained over augmented subsets obtains a more accurate estimate of the plume shape and location.

2) *PDE-Informed Neural Networks*: The physics-informed neural network (PINN [19], [36]) is currently the most widely used approach to building PDE-informed neural networks. Various other PDE-informed

neural network frameworks have also been applied to seismic FWI [37], [38], [39], [40]. All PDE-informed neural networks aim to integrate the physics into the trained network in some way. The most widely used approach is to augment the loss function with a residual term from the governing equation, which regularizes the solution space during training. The main benefits to incorporating the wave equation have been found to be improved network accuracy, generalization and robustness, reduction of the amount of training data required, and in some cases, removal of the need for labeled data. In the following, we discuss the method with which we are most familiar. For example, Jin et al. [38] developed an Unsupervised Physical-informed Full-Waveform Inversion (UPFWI), which can be posed as

$$\widehat{\mathbf{m}}(\mathbf{d}) = g_{\theta^*}(\mathbf{d}), \text{ where } \theta^*(\Phi_u) = \arg \min_{\theta} \sum_{\mathbf{d}_i \in \Phi_u} \mathcal{L}(f(g_{\theta}(\mathbf{d}_i)), \mathbf{d}_i), \quad (6)$$

where  $\Phi_u$  represents an unsupervised dataset that contains seismic data  $\mathbf{d}_i$  alone. Finite differences were used to approximate the forward modeling of PDE as a differentiable operator (from velocity map to seismic data), and then model its inversion by CNN (from seismic data to velocity map). Hence, the supervised data-driven inversion task in Eq. (5) is transformed into the unsupervised seismic data reconstruction task in Eq. (6), where the governing wave equation plays an essential role. Comparing the network architecture of UPFWI (Fig. 6) to InversionNet (Fig. 2), the major difference is the explicit incorporation of wave equation. It constrains the output of the CNN (the orange network in Fig. 6), when passing through an explicit wave equation, to match the original input (i.e., seismic data). Some additional physical constraints are imposed to further enhance the imaging results: one is to leverage a learning-based perceptual loss function by extracting style and content features of the output seismic data via VGG16 network, and the other is to confine the resulting dimension of the CNN output, which should be the same as that of the expected velocity maps. The former helps in imaging the reflection layers of the velocity maps while the latter is helpful to some degree to constrain the solution space.

We provide a performance comparison of UPFWI with two pure data-driven methods (i.e., InversionNet [7] and VelocityGAN [41]), and a supervised physics-informed method H-PGNN [37] using two different datasets (FlatFault and CurvedFault [42]). UPFWI performs comparably to all three baseline methods on both FlatFault and CurvedFault. The velocity maps inverted by different methods are shown in Figure 7, which is consistent with the quantitative analysis in Table III.

3) *Other Physics-Informed Constraints:* Another interesting way of incorporating physics is to estimate the velocity maps based on a learned representation of the underlying distribution of the physical parameters. Specifically, the generative adversarial network (GAN) is a popular network architecture that has been used to characterize data distribution. A number of recent works are representative of this

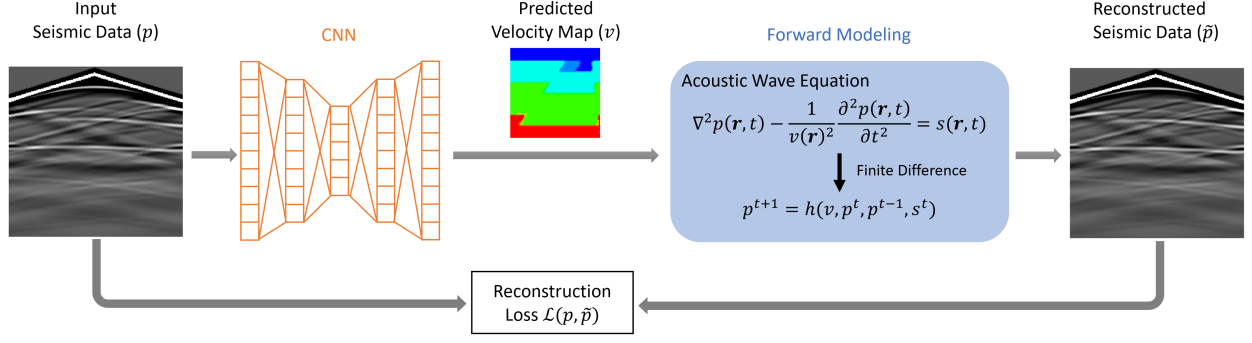


Fig. 6: Schematic illustration of UPFWI structure, which comprises a CNN to learn an inverse mapping and a differentiable operator to approximate the forward modeling of PDE [38].

Supervision	Method	FlatFault			CurvedFault		
		MAE ↓	MSE ↓	SSIM ↑	MAE ↓	MSE ↓	SSIM ↑
Supervised	InversionNet [7]	15.83	2156.00	0.9832	23.77	5285.38	0.9681
	VelocityGAN [41]	16.15	1770.31	0.9857	25.83	5076.79	0.9699
	H-PGNN+ [37]	<b>12.91</b>	1565.02	0.9874	24.19	5139.60	0.9685
Unsupervised	UPFWI	14.60	<b>1146.09</b>	<b>0.9895</b>	<b>23.56</b>	<b>3639.96</b>	<b>0.9756</b>

TABLE III: Quantitative results evaluated in MAE, MSE and SSIM. UPFWI yields better inversion accuracy comparable to that of supervised baselines.

category [43], [41], [44]. Mathematically, these ideas can be formulated as:

$$\begin{aligned}
 \widehat{\mathbf{m}}(\mathbf{z}^*) &= g_{\theta^*}(\mathbf{z}^*) \\
 \text{s.t. } \mathbf{z}^*(\mathbf{d}) &= \arg \min_{\mathbf{z}} \mathcal{L}(f(g_{\theta^*}(\mathbf{z})), \mathbf{d}) \\
 \theta^*(\Phi_m) &= \arg \min_{\theta} \sum_{\mathbf{m}_i \in \Phi_m} \mathcal{L}_{\text{GAN}}(g_{\theta}(\alpha_i), \mathbf{m}_i),
 \end{aligned} \tag{7}$$

where  $\Phi_m$  is a training dataset including numerous velocity maps; and  $\mathbf{z}$  and  $\alpha_i$  are tensors sampled from the normal distribution. The iterative optimization is then performed on  $\mathbf{z}$  to draw a velocity map sampled from the adaptive distribution. It is worth mentioning a major difference in learning strategy comparing Eq. (7) with Eq. (6). The optimal weights  $\theta^*(\Phi_u)$  in Eq. (6) is obtained based on the dataset  $\Phi_u$  for a subdomain, whereas the optimal  $\theta^*(\mathbf{d})$  in Eq. (7) is optimized per sample  $\mathbf{d}$ . After training with sufficient seismic data, Eq. (6) yields a mapping from seismic data to velocity maps that can be directly used for inference without any optimization. In contrast, the optimal weights in Eq. (7) are sample dependent, requiring optimization for any new sample before inference.

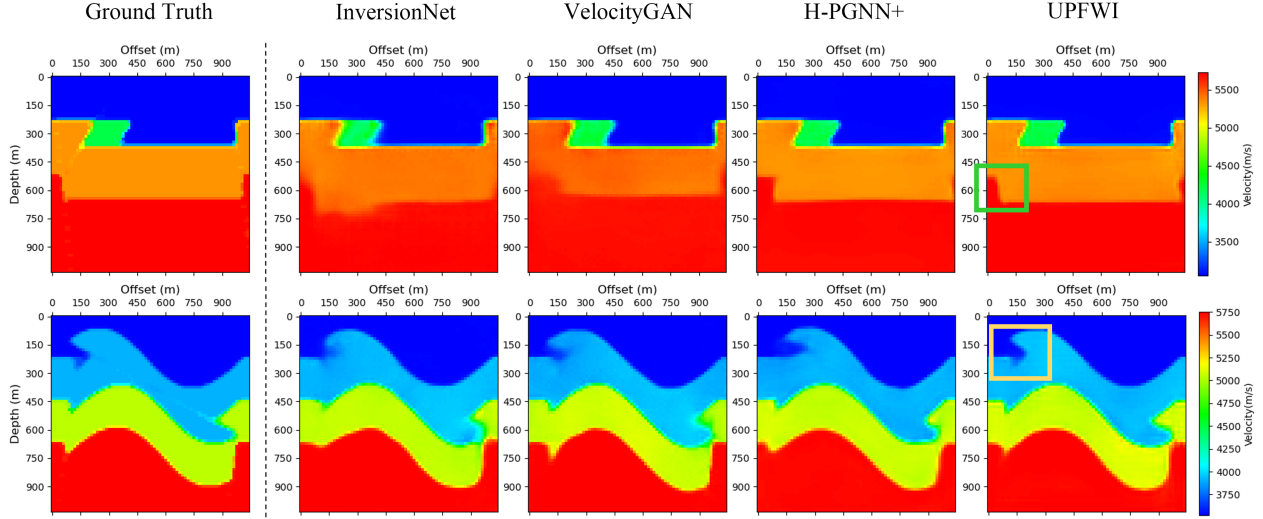


Fig. 7: Comparison of different methods on inverted velocity maps of FlatFault (top) and CurvedFault (bottom). For FlatFault, UPFWI reveals more accurate details at layer boundaries and the slope of the fault in deep region. For CurvedFault, UPFWI reconstructs the geological anomalies on the surface that best match the ground truth.

### C. Physics Models Enhanced by Machine Learning

Three groups of methods are included in this category, as shown in the green hexagons in Fig. 3. These methods are built on physics-based iterative approaches, with machine learning techniques introduced to enhance the computational techniques. Specifically, we will discuss how to learn useful prior knowledge using deep learning models for physics-based methods, how to leverage ML to estimate gradient, and how to learn an effective loss function using ML for physics-based methods.

1) *Prior Knowledge Learning*: A large body of work has demonstrated the importance of prior knowledge in improving the performance of physics-based FWI approaches. This prior knowledge includes the existing parameter distribution [45], [46], [47], [48], the low-frequency component of the measured signals [49], [50], [51], and others. Here our focus is on exploiting the known parameter distribution, which has been proposed for a number of data-driven methods. These techniques originate from the general idea of deep internal learning that trains deep neural networks on a single training sample. An example would be the deep image prior [52] in which deep learning models are utilized to learn useful prior knowledge via re-parameterization. When applying similar ideas to seismic inversion problems, it has already been shown that a couple of the major benefits are expected including: one, the techniques mitigate local minima issue due to the sparse representation of the velocity maps using re-parameterization; two, the re-parameterization can be seen as an implicit regularization, which in turn helps alleviate the ill-posedness






Fig. 8: Schematic illustration of the idea of reparameterization, which combines neural network and PDE of seismic waves. The neural network is used to parameterize the velocity maps while the PDE is to calculate the forward modeling and data misfit term (Image courtesy of [45] and used with permission.)

of the inversion problem [47]. Schematically, the network structures of those methods usually consist of two components: a neural network and a PDE, which is illustrated in Fig. 8. The neural network is used to re-parameterize the velocity maps while the PDE is to calculate the forward modeling procedure and obtain the data misfit. Mathematically, seismic inversion with re-parameterization can be formulated as

$$\begin{aligned} \widehat{\mathbf{m}}(\mathbf{d}) &= g_{\theta^*(\mathbf{d})}(\boldsymbol{\alpha}) \\ \text{s.t. } \theta^*(\mathbf{d}) &= \arg \min_{\boldsymbol{\theta}} \mathcal{L}(f(g_{\boldsymbol{\theta}}(\boldsymbol{\alpha})), \mathbf{d}), \end{aligned} \quad (8)$$

where  $\boldsymbol{\alpha}$  is a random tensor drawn from a normal distribution, i.e.,  $\boldsymbol{\alpha} \sim \mathcal{N}(0, \sigma^2)$ . Different network architectures,  $g_{\boldsymbol{\theta}}(\boldsymbol{\alpha})$ , have been explored to study their regularization implicit effect in those existing work including CNN-domain FWI [46], DNN-FWI [47] and FCN-FWI [45]. As shown in Eq. (8), these methods involve iterative optimization, with  $\theta^*(\mathbf{d})$  in Eq. (8) being optimized for seismic measurements  $\mathbf{d}$  for a particular time and site, and requiring re-training in order to compute the inversion for a different input  $\mathbf{d}$ . Unlike those subdomain approaches, the optimal weights  $\theta^*(\Phi_u)$  in Eq. (5) are obtained based on the dataset  $\Phi_u$ , thus one network would be applicable to multiple seismic samples once fully trained. Here, the training of the networks,  $g_{\boldsymbol{\theta}}(\boldsymbol{\alpha})$ , is usually carried out in two steps. The networks are first pre-trained on an expert initial guess of the solution, which can be obtained, for instance, via traveltime inversion techniques or migration velocity analysis. After pre-training, the network is used to reparameterize the velocity and the network parameters and is updated iteratively via the computation of the gradient. Comparable results to physics-based inversion methods have been reported in the aforementioned works. There are, however, some technical challenges associated with these methods. One is the heavy reliance on the initial guess of the solution. Eq. (8) requires an *expert* initial guess, which is crucial for optimization

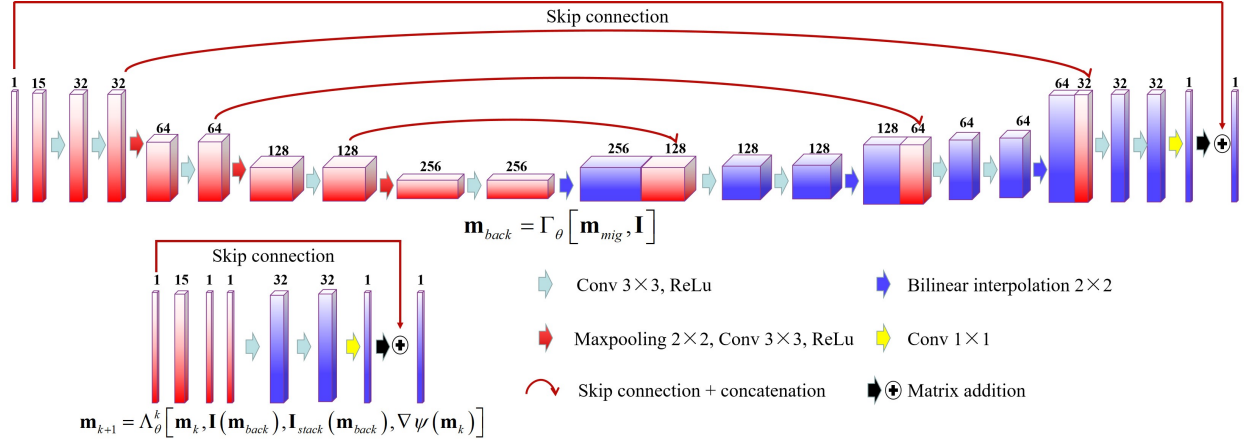


Fig. 9: Schematic illustration of the idea of (Image courtesy of [53] and used with permission.)

and expensive to acquire [47]. An unsatisfactory initial guess would potentially mislead the training and yield a completely wrong inversion. The second is the high computational cost in training, which is due to the nature of the per sample learning strategy. The optimal network parameters,  $\theta$  in Eq. (8), will not be able to be generalized to other data without a full re-training.

2) *Unrolling-Based Techniques*: The computational techniques solving Eq. (3) are based on iterative methods. The main idea behind this group of methods is to unroll an iterative schemes with a small number of iterations, replacing critical components of the scheme, such as the computation of the gradient of the data fidelity term, with a trainable module. Mathematically, this can be represented as

$$\mathbf{m}_{k+1} = \mathbf{m}_k + \alpha_k \nabla E(\mathbf{m}_k) \approx \Psi_{\theta}^k[\mathbf{m}_k, \nabla E(\mathbf{m}_k)], \quad k = 0, 1, \dots, N-1, \quad (9)$$

where  $\alpha_k$  is the line search step size at each iteration;  $\nabla E(\mathbf{m}_k)$  is the gradient term; and  $\Psi_{\theta}^k$  is a neural network with learnable parameters  $\theta$ . With an appropriately trained network of  $\Psi_{\theta}^k$ , the gradient of FWI will be obtained efficiently. However, there are two major issues in employing Eq. (9). One is the irreversible loss of some information when calculating the gradient based on the cross-correlation of the forward and backward propagated wavefields. A heuristic interpretation is that the original raw seismic waveform data are mapped from high-dimensional space into a low-dimensional space where the gradient resides in, resulting in unavoidable lost of some information of the observed data. Two is that the accuracy of Eq. (9) depends on a good initial guess of the solution. A variety of approaches have been developed to address these issues [53], [54], [55], [56]. Specifically, Zhang and Gao [53], [54] developed a two-stage learning strategy by unrolling the FWI iteration. More recently, the unrolling technique was combined with forward modeling and prior knowledge [56]. The proposed approach involved nested iterations, with the inner loop corresponding to prediction of the velocity map from the training set and

migration image, and the outer loop improving the training set using the velocity map estimated within the inner loop. Here, we provide details of the two-stage unrolling technique [53], which is illustrated in Fig. 9.

- Stage 1 – a network is designed to estimate background velocity map  $\mathbf{m}_{\text{back}}$  from input

$$\mathbf{m}_{\text{back}} = \Gamma_{\theta} [\mathbf{m}_{\text{smooth}}, I(\mathbf{m}_{\text{smooth}})] , \quad (10)$$

where  $\Gamma_{\theta}$  is an encoder-decoder neural network with learnable parameter  $\theta$  (shown in the upper panel of Fig. 9);  $\mathbf{m}_{\text{back}}$  is the reconstructed background velocity map;  $\mathbf{m}_{\text{smooth}}$  is the smoothed background velocity map; and  $I(\cdot)$  is the reverse time migration operator<sup>6</sup> using prestack seismic data<sup>7</sup>.

- Stage 2 – the high resolution velocity map is updated iteratively from the background velocity map as

$$\mathbf{m}_{k+1} = \Lambda_{\theta}^k [\mathbf{m}_k, I(\mathbf{m}_{\text{back}}), I_{\text{stack}}(\mathbf{m}_{\text{back}}), \nabla R(\mathbf{m}_k)] , \quad k = 0, 1, \dots, N-1, \quad (11)$$

where  $\Lambda_{\theta}^k$  is a residual neural network with learnable parameter  $\theta$  (shown in the lower panel of Fig. 9);  $I_{\text{stack}}(\cdot)$  is the reverse time migration operator using the post-stacked seismic data<sup>8</sup>; and  $\nabla R(\mathbf{m}_k)$  is the gradient of the regularization in Eq. (3).

An advantage of the two-stage learning strategy (Eqs. (10) and (11)) over the standard unrolling method as in Eq. (9) is that the neural network input and output are in the same space. In addition, more physical information (i.e., reverse time migration images) are used, which not only provides more details to reconstruct the velocity maps, but also alleviates the model reliance on a good initial guess of the solution. As mentioned in [53], the prediction performance of the model heavily relies on the quality and quantity of the samples and labels. Some performance improvements are reported in the literature due to the incorporation of additional physical information, but this is more computationally expensive because of the nature of the iterative procedure.

*3) Learning Loss Functions:* The loss function plays a critical role in physics-based inversion techniques. An inappropriately selected loss function will make the iteration prone to non-convexity/multiple minima problems and over-sensitivity to the initial guess. Traditional methods for tackling this challenge

<sup>6</sup>Reverse time migration is a seismic processing technique that is able to estimate limited subsurface structural information, such as the locations of layer boundaries.

<sup>7</sup>Prestacking is a seismic data pre-processing technique and prestack data is the data before stack/migration.

<sup>8</sup>Post-stack data refers to the data that is processed after stacking and results in a subsurface structural image, which is in the time domain. Here, the stacking operation is the process to sum seismic traces to improve the signal-to-noise ratio.

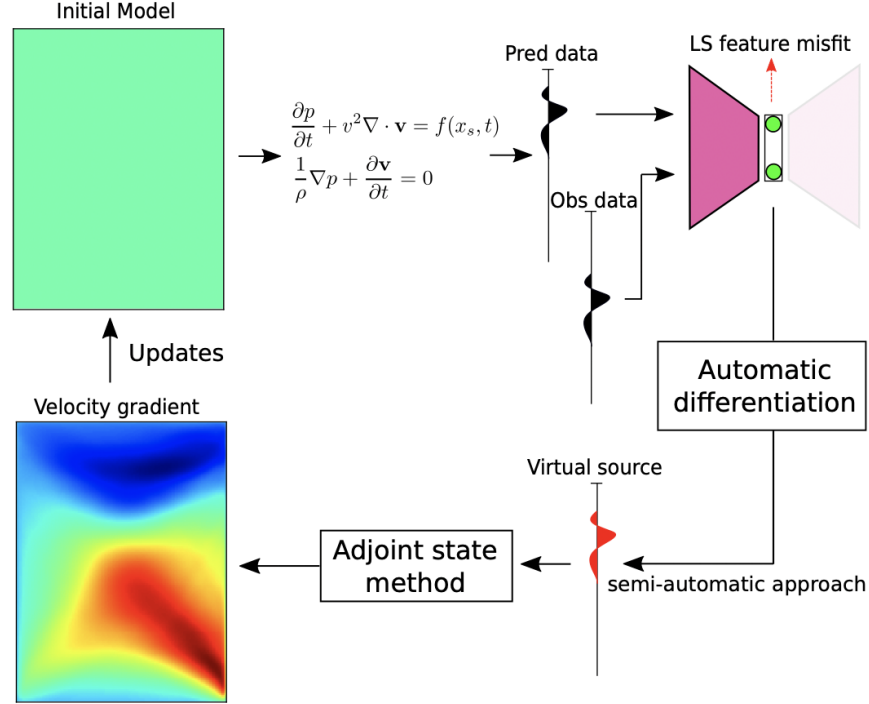


Fig. 10: Schematic illustration of the hybrid machine learning inversion framework (Image courtesy of [57] and used with permission).

are based on hand-crafted functions, which may succeed in some datasets but fail in others [3]. To overcome this issue, ML approaches have been utilized to learn a data fidelity function [57], [58], [59]. Specifically, a learning-based data fidelity function that measures high-level features has been shown to be effective in mitigating the local minima issue. Following this, Chen and Saygin [57] develop a multi-scale hybrid machine learning inversion method by defining a data fidelity function that measures the latent space features of the observed and predicted data (see Fig. 10). The velocity maps are updated iteratively by minimizing the learning-based data fidelity term. A different perspective is to use a network to approximate either the data fidelity function or the regularization term. Sun and Alkhalifah [58] employ a meta-learning methodology to learn a robust data fidelity function for FWI. Zhang and Alkhalifah [59] trained a deep neural network to first estimate the distribution of facies (the overall characteristics of a rock unit that reflect its origin and differentiate the unit from others around it), which will be used as regularization to constrain the iterative procedure of the FWI. We select the work in [57] to represent this group of approaches due to its promising performance and applicability to real data. Their idea is illustrated in Fig. 10. Rather than using the traditional loss term (i.e.,  $\|d - f(\mathbf{m})\|_2^2$  in Eq. (3)), a new




Fig. 11: Overview of the datasets in OpenFWI [42].

loss function is proposed

$$E(\mathbf{m}) = \min_{\mathbf{m}} \left\{ \|W_{\theta}(\mathbf{d}) - W_{\theta}(f(\mathbf{m}))\|_2^2 \right\}, \quad (12)$$

where  $W_{\theta}$  is a neural network with learnable parameter  $\theta$  used to extract useful latent space features from both real data,  $\mathbf{d}$  and predicted data  $f(\mathbf{m})$ . The gradient of the network can be obtained via automatic differentiation and the adjoint-state method. The proposed learning-based loss function in Eq. (12) is found to be less prone to entrapment in local minima compared to the traditional loss in Eq.(3). Furthermore, the proposed inversion method can recover the subsurface velocity structure in different scales. A low dimensional latent space contains the kinematic information of the seismic waveforms, while large-dimensional latent space features preserve the dynamic information from the waveforms. Experiments with both synthetic and real data tests have been carried out to demonstrate the performance of the proposed method [57].

#### D. Open Data, Benchmarks, and Software

1) *Open Data and Benchmarks*: The availability of an extensive database is vital for machine learning-based inversion techniques. Unlike some data-rich domains (such as computer vision), this is problematic for seismic inversion in two ways. One is that data are often proprietary and hence challenging to share, but the second and more important problem is that acquiring correctly labeled data is extremely difficult or even impossible as the true velocity maps are never known exactly. There have been some recent effort to address this urgent need for more and better open data for the R&D of seismic inversion. The U.S.

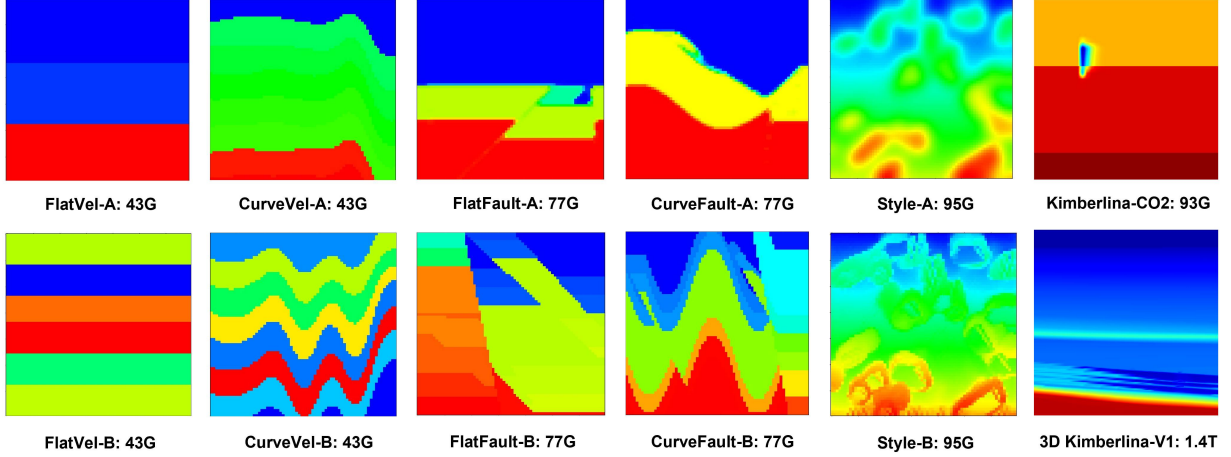


Fig. 12: Gallery of OPENFWI, which contains one example of velocity maps from each dataset in OPENFWI.

Department of Energy (DOE) created the Energy Data eXchange (EDX), which is a virtual library and data laboratory built to find, connect, curate, use and re-use data to advance fossil energy and environmental R&D [60]. EDX supports the entire life cycle of data by offering secure, private collaborative workspaces for ongoing research projects until they mature and become catalogued, curated, and published. Recently, Los Alamos National Laboratory created and released OPENFWI<sup>9</sup>, a collection of large-scale multi-structural benchmark datasets, to facilitate diversified, rigorous, and reproducible research on FWI [42]. As shown in Fig. 11, OPENFWI consists of 12 datasets (2.1TB in total) synthesized from multiple sources. The dataset encompasses diverse domains in geophysics (interface, fault, CO<sub>2</sub> reservoir, etc.), covers different geological subsurface structures (flat, curve, etc.), and contains various amounts of data samples (2K - 67K). It also includes a dataset for 3D FWI. We provide in Fig. 12 an example of velocity maps from each dataset. Along with the dataset, an extensive study on performance benchmarks using both physics-based inversion and data-driven techniques are provided. In addition to evaluations on a single dataset, OPENFWI enables the study of generalization across datasets.

2) *Software*: Several software packages, including NVIDIA Modulus [61], DeepXDE [62], and PyDens [63] were released since 2019. Instead of targeting at FWI problems, those packages are mostly designed for a general-purpose PINN framework for various applications. Within the scope of the interests

<sup>9</sup>OpenFWI: Datasets accessible at <https://openfwi-lanl.github.io/>.

Method	Learning Strategy		Supervision		Physics Knowledge				
	Data-Driven	Hybrid	Supervised	Unsupervised	Augmentation	Governing Eq.	Priors	Unrolling	Loss
InversionNet [7]	✓		✓						
Araya-Polo et al. [8]	✓		✓						
ModifiedFCN [9]	✓		✓						
Li et al. [6]	✓		✓						
Renan et al. [33]		✓	✓		✓	✓			
Yang et al. [31]		✓	✓		✓			✓	
Feng et al. [32]		✓	✓		✓			✓	
VelocityGAN [41]		✓	✓					✓	
UPFWI [38]		✓		✓		✓			✓
H-PGNN+ [37]		✓	✓			✓			
NNFWI [45]		✓		✓		✓			
He & Wang [47]		✓		✓		✓			
Wu & McMechan [46]		✓		✓		✓			
Dhara and Sen [40]		✓	✓			✓			
Rasht-Behesht et al. [36]		✓	✓			✓			
Mosser et al. [43]		✓	✓					✓	
Richardson et al. [44]		✓	✓					✓	
Zhang & Gao [53]		✓	✓						✓
Zhang et al. [55]		✓	✓						✓
CNN-RWI [56]		✓	✓			✓	✓	✓	
Chen & Saygin [57]		✓		✓					✓
Zhang & Alkhalifah [59]		✓		✓					✓
Sun & Alkhalifah [58]		✓		✓					✓

TABLE IV: Summary of machine learning based seismic inversion techniques addressed in this paper.

of this paper, there are some open packages: InversionNet [7], VelocityGAN [41]<sup>10</sup>, and CNN-RWI<sup>11</sup> [56]. Both InversionNet and VelocityGAN are PyTorch-based packages for solving FWI based on acoustic wave equation. Extensive examples, tests, and benchmarks are provided for running both codes on OpenFWI datasets. A thorough tutorial is also provided to help with using codes on the OpenFWI website.

<sup>10</sup>InversionNet and VelocityGAN: Both were recently released by Los Alamos National Laboratory. Source codes, examples, and tutorial are accessible at <https://github.com/lanl/OpenFWI>.

<sup>11</sup>CNN-RWI: Souce codes are accessible at <https://github.com/YulangWu/CNN-RWI>.

#### IV. CHALLENGES AND FUTURE RESEARCH OPPORTUNITIES

A summary of the inversion techniques mentioned in this review is provided in Table IV. This collection of papers is by no means exhaustive, but it does represent some of the major R&D efforts in the past five years. While physics-informed data-driven techniques have already shown great potential in addressing some of the existing issues in seismic inversion (such as high computational cost, non-unique solutions, etc), new challenges have been encountered during the development of those techniques. Here we briefly discuss some of these, and propose potential future research topics.

- 1) **Generalization and Model Robustness.** These two topics are critical to any machine learning problems, but especially to seismic imaging. Generalization formally refers to the issue of performance on out-of-sample data (data that is not in the training set), but for seismic imaging, the concern extends to data that is drawn “out of distribution,” or from a drifted distribution. Robustness refers to the performance degradation of a model with perturbed noisy test data. Incorporation of physics seeks to enhance both the generalization and the robustness, though it is difficult to explicitly demonstrate the effectiveness of the models on field data, which are usually contaminated with a significant amount of noise and artifacts. We remark that Deng et al. [64] used theoretical arguments to provide a rigorous derivation of worst-case error bounds in seismic imaging.
- 2) **Scientific “Sim2Real.”** Given that data are scarce in subsurface geoscience, full-physics simulations provide some compensation for this lack of data. How to bridge the “knowledge gap” (the gap between physical simulations and real applications) by utilizing machine learning and the underlying physics is a scientifically important question that needs to be carefully addressed. A similar challenge also exists in other scientific disciplines (such as robotics), where the notion of *Sim2Real* has been proposed to transfer knowledge learned in simulation to the real data [65]. That leads to two potential solutions: one is the pure model-based simulators and the other is machine learning methods (such as generative models). The former usually simplifies the complex physics resulting in an unavoidable reality gap between the simulation and real data. The latter may be used as a compensating tool to learn underlying physics directly out of data. We are beginning to see new research and development has been carried out in both areas. However, the feasibility of these approaches on real applications with field data needs to be validated.
- 3) **Penalties and Constraints.** For most of the methods in Table IV, the physics knowledge is embedded in the loss function as soft constraints via penalties such as regularization terms. Mathematically, these problems can be generally posed as

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} \{ \mathcal{L}(g_{\boldsymbol{\theta}}(\mathbf{m}), \mathbf{d}) + \lambda \mathcal{R}(f(\mathbf{d}), \mathbf{m}) \} , \quad (13)$$



where  $\mathbf{m}$  and  $\mathbf{d}$  are the inputs and predictions, respectively;  $g_{\theta}(\cdot)$  represents the mapping relationship fitted by the network model with learnable parameters  $\theta$ ;  $f(\cdot)$  is the mathematical representation of the domain knowledge (such as the governing PDEs); and  $\lambda$  is the regularization parameter connecting the loss term  $\mathcal{L}(\cdot)$  and the regularization term  $\mathcal{R}(\cdot)$ . With soft constraint as posed in Eq. (13), the physics knowledge can be only weakly enforced in the prediction results, meaning that they can not ensure the predictions do not violate the underlying physics constraints. To overcome this limitation of soft constrained approaches, different optimization techniques are recently proposed to incorporate physics knowledge as hard constraints, which ensures a given governing equation is strictly satisfied in a given domain [66].

- 4) **Uncertainty Quantification.** Considering issues of generalization, model robustness, and information gap between synthetic and field data, it is important to quantitatively evaluate the model performance on both inversion accuracy and the associated confidence. Uncertainty quantification (UQ) has been widely used as a tool to achieve this. There has been some recent work studying data-driven seismic inversion with UQ to provide insights into the various UQ induced by different perturbations including measurement noise and model parameters [67], [45].
- 5) **Lightweight Networks.** State-of-the-art ML-based FWI solvers employ encoder-decoder CNNs to learn a direct mapping from waveform measurements to physical properties. Usually, CNNs with multiple layers of convolutional blocks are built to learn sufficient features from local patches to global image regions. This is a reasonable approach and works for many inverse problems. However, it ignores one critical fact: The measurements and properties are more than just 2D images. They both have physical meanings and are connected by the governing wave equation. Here, an open question is how to leverage physics information to not only reduce the network complexity (depth and width) but also maintain a good inversion accuracy. There have been some interesting methods targeting this direction [68].
- 6) **3D Full Waveform Inversion.** In reality, of course, the subsurface is a 3D volume, not a 2D plane. Almost all the literature that we find are developed for 2D problems, which limit their applications in real-world scenarios. However, extending from 2D to 3D inversion is not straightforward. Some major technical barriers need to be carefully addressed for real 3D FWI methods. It includes high computational complexity, overfitting risk, and large memory consumption. Particularly, 3D convolutions on input, output and intermediate activations in middle network layers are computationally intensive. Large-scale floating-point operations will increase the time usage in both training and testing. Furthermore, a conventional 3D CNN consists of tens of millions of parameters, which leads to a much higher overfitting risk. Training a normal 3D CNN with the minimum batch size

of 1 occupies nearly all the available memory on a graphics processing unit. Moreover, training with excessively small batch sizes may degrade the imaging performance. A most recent technique on ML-based 3D FWI have been developed [69]. More research effort would be needed along this line.

- 7) **Joint Multi-geophysical Inversion.** Complex subsurface systems usually involve multiple geophysical exploration techniques. The issue of coupling models of different geophysical signals at different scales and governed by different physical laws is wide open and represents an enormously challenging area for various subsurface applications. One particular topic of great interest is to extract small but critical signatures from a noisy environment, which will have an immediate impact on rare anomaly detection and complex geophysical system monitoring. A couple of recent works combining seismic waveform inversion and electromagnetic inversion have shown great promise in monitoring CO<sub>2</sub> migration [70], [71].
- 8) **Realistic Applications.** Applying physics-based FWI methods to real seismic data is challenging and even more so for emerging ML-based methods. While there are some notable papers demonstrating the application of ML-based FWI methods to real data [57], [32], [72], most of the existing work still utilizes synthetic data. ML-based approaches share some of the challenges of physics-based FWI methods, and often introduce some of their own. Most existing ML approaches are based on simplified forward model built on the acoustic wave equation, while real seismic wave propagation is somewhat more complex, with properties such as anisotropy and visco-elasticity. Environmental noise and data acquisition artifacts are also an important factors for the successful implementation of FWI methods, yet most current ML approaches only consider simplified synthetic noise such as Gaussian noise. Thus, there is room to improve ML approaches by incorporating more sophisticated noise and data artifact models, as well as more realistic wave physics.

## V. CONCLUSIONS

The problem of seismic full-waveform inversion, first formulated more than 35 years ago, continues to be a major area of research in subsurface geoscience. From subsurface velocity building for fossil-fuel reservoirs in subsurface energy exploration to imaging global subduction zones in earthquake seismology, there is a demand for algorithms that are *faster*, *more accurate*, and *scalable* to very large datasets. As demonstrated by the research efforts highlighted in this review, scientific machine learning is showing great potential on all of these fronts. Serious challenges remain – direct subsurface measurements are difficult and necessarily sparse; generalizing learned solutions to regimes outside their training is inherently fraught with uncertainty; networks should be robust to inevitable measurement and noise artifacts – but we believe

that there are ample opportunities for physics-informed machine learning to enable exciting progress in this science-rich and data-starved field.

## ACKNOWLEDGMENT

This work was funded by the Los Alamos National Laboratory (LANL) - Laboratory Directed Research and Development program under project numbers 20210542MFR and 20200061DR, and by the U.S. Department of Energy (DOE) Office of Fossil Energy’s Carbon Storage Research Program via the Science-Informed Machine Learning to Accelerate Real Time Decision Making for Carbon Storage (SMART-CS) Initiative.

## REFERENCES

- [1] Albert Tarantola. Inverse problem theory and methods for model parameter estimation. SIAM, 2005.
- [2] Albert Tarantola. Inversion of seismic reflection data in the acoustic approximation. Geophysics, 49(8):1259–1266, 1984.
- [3] Jean Virieux and Stephane Operto. An overview of full-waveform inversion in exploration geophysics. Geophysics, 74(6):WCC1–WCC26, 2009.
- [4] Amir Adler, Mauricio Araya-Polo, and Tomaso Poggio. Deep learning for seismic inverse problems: Toward the acceleration of geophysical analysis workflows. IEEE Signal Processing Magazine, 38:89–119, 2021.
- [5] Jeroen Tromp. Seismic wavefield imaging of earth’s interior across scales. Nature Review Earth and Environment, 1:40–53, 2020.
- [6] Shucai Li, Bin Liu, Yuxiao Ren, Yangkang Chen, Senlin Yang, Yunhai Wang, and Peng Jiang. Deep-Learning Inversion of Seismic Data. IEEE Transactions on Geoscience and Remote Sensing, 58(3):2135–2149, March 2020.
- [7] Yue Wu and Youzuo Lin. InversionNet: An efficient and accurate data-driven full waveform inversion. IEEE Transactions on Computational Imaging, 6:419–433, 2019.
- [8] Mauricio Araya-Polo, Joseph Jennings, Amir Adler, and Taylor Dahlke. Deep-learning tomography. The Leading Edge, 37(1):58–66, 2018.
- [9] Wenlong Wang, Fangshu Yang, and Jianwei Ma. Velocity model building with a modified fully convolutional network. In SEG, pages 2086–2090, 2018.
- [10] Òscar Agudo. Acoustic full-waveform inversion in geophysical and medical imaging. PhD thesis, Imperial College London, 2018.
- [11] Karianne Bergen, Paul Johnson, Maarten V. de Hoop, and Gregory Beroza. Machine learning for data-driven discovery in solid earth geoscience. Science, 363, 2019.
- [12] Gregory Ongie, Ajil Jalal, Christopher Metzler, Richard Baraniuk, Alexandros Dimakis, and Rebecca Willett. Deep learning techniques for inverse problems in imaging. IEEE Journal on Selected Areas in Information Theory, 1:39 – 56, 2020.
- [13] Xiwei Yu and Jianwei Ma. Deep learning for geophysics: Current and future trends. Reviews of Geophysics, 59:e2021RG000742, 2021.
- [14] Jared Willard, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. Integrating physics-based modeling with machine learning: A survey. arXiv preprint arXiv:2003.04919, 2020.
- [15] George Karniadakis, Ioannis Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. Nature Reviews Physics, 3:422–440, 2021.

- [16] Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli. Scientific machine learning through physics-informed neural networks: Where we are and what's next. Journal of Scientific Computing, 92, 2022.
- [17] Lluís Guasch, Oscar Calderón, Meng-Xing Tang, Parashkev Nachev, and Michael Warner. Full-waveform inversion imaging of the human brain. NPJ Digital Medicine, 3(28), 2020.
- [18] Kun Wang, Thomas Matthews, Fatima Anis, Cuiping Li, Neb Duric, and Mark Anastasio. Waveform inversion with source encoding for breast sound speed reconstruction in ultrasound computed tomography. IEEE Trans Ultrason Ferroelectr Freq Control, 62(3):475–493, 2015.
- [19] Maziar Raissi, Paris Perdikaris, and George Karniadakis. 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, 2019.
- [20] Jared Willard, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. Integrating scientific knowledge with machine learning for engineering and environmental systems. arXiv, 2003.04919, 2020.
- [21] Youzuo Lin and Lianjie Huang. Acoustic- and elastic-waveform inversion using a modified Total-Variation regularization scheme. Geophysical Journal International, 200:489–502, 2015.
- [22] Yong Ma and Dave Hale. Wave-equation reflection traveltime inversion with dynamic warping and full-waveform inversion. Geophysics, 78(6):R223–R233, 2013.
- [23] Zhigang Zhang and Lianjie Huang. Double-difference elastic-waveform inversion with prior information for time-lapse monitoring. Geophysics, 78(6):R259–R273, 2013.
- [24] Khiem Tran, Michael McVay, Michael Faraone, and David Horhota. Sinkhole detection using 2D full seismic waveform tomography. Geophysics, 78(5):R175–R183, 2013.
- [25] Antoine Guitton, Gboyega Ayeni, and Esteban Díaz. Constrained full-waveform inversion by model reparameterization. Geophysics, 77(2):R117–R127, 2012.
- [26] Peyman Moghaddam and Felix Herrmann. Randomized full-waveform inversion: A dimensionality-reduction approach. In SEG, pages 4453–4457, 2010.
- [27] Leonardo Azevedo. Model reduction in geostatistical seismic inversion with functional data analysis. Geophysics, 87:1JF–WA102, 2022.
- [28] Shihang Feng, Xitong Zhang, Brendt Wohlberg, Neill P. Symons, and Youzuo Lin. Connect the dots: In situ 4-D seismic monitoring of CO<sub>2</sub> storage with spatio-temporal CNNs. IEEE Transactions on Geoscience and Remote Sensing, 60:1–16, 2022.
- [29] Zhaoqi Gao, Chuang Li, Bing Zhang, Xiudi Jiang, Zhibin Pan, Jinghuai Gao, and Zongben Xu. Building large-scale density model via a deep-learning-based data-driven method. Geophysics, 86:1JF – V89, 2021.
- [30] Jonathan Downton, Olivia Collet, and Daniel Hampson. Theory-guided data science-based reservoir prediction of a north sea oil field. The Leading Edge, 39:689 – 764, 2020.
- [31] Yuxin Yang, Xitong Zhang, Qiang Guan, and Youzuo Lin. Making invisible visible: Data-driven seismic inversion with spatio-temporally constrained data augmentation. IEEE Transactions on Geoscience and Remote Sensing, 2022 (Accepted).
- [32] Shihang Feng, Youzuo Lin, and Brendt Wohlberg. Multiscale data-driven seismic full-waveform inversion with field data study. IEEE Transactions on Geoscience and Remote Sensing, pages 1–14, 2021.
- [33] Renán Rojas-Gómez, Jihyun Yang, Youzuo Lin, James Theiler, and Brendt Wohlberg. Physics-consistent data-driven waveform inversion with adaptive data augmentation. IEEE Geoscience and Remote Sensing Letters, 19:1–5, 2022.

- [34] Tariq Alkhalifah, Hanchen Wang, and Oleg Ovcharenko. MLReal: Bridging the gap between training on synthetic data and real data applications in machine learning. arXiv, 2021.
- [35] Preston Jordan and Jeffrey Wagoner. Characterizing construction of existing wells to a CO<sub>2</sub> storage target: The Kimberlina site, California. Technical report, U.S. Department of Energy - Office of Fossil Energy, March 2017.
- [36] Majid Rasht-Behesht, Christian Huber, Khemraj Shukla, and George Em Karniadakis. Physics-informed neural networks (PINNs) for wave propagation and full waveform inversions. Journal of Geophysical Research: Solid Earth, 127:e2021JB023120, 2022.
- [37] Jian Sun, Kristopher A Innanen, and Chao Huang. Physics-guided deep learning for seismic inversion with hybrid training and uncertainty analysis. Geophysics, 86(3):R303–R317, 2021.
- [38] Peng Jin, Xitong Zhang, Yinpeng Chen, Sharon Huang, Zicheng Liu, and Youzuo Lin. Unsupervised learning of full-waveform inversion: Connecting CNN and partial differential equation in a loop. In The Tenth International Conference on Learning Representations (ICLR), 2022.
- [39] Chao Song, Tariq Alkhalifah, and Umair Bin Waheed. Solving the frequency-domain acoustic VTI wave equation using physics-informed neural networks. Geophysical Journal International, 225:846 – 859, 2021.
- [40] Arnab Dhara and Mrinal Sen. Physics-guided deep autoencoder to overcome the need for a starting model in full-waveform inversion. The Leading Edge, 41:375–381, 2022.
- [41] Zhongping Zhang and Youzuo Lin. Data-driven seismic waveform inversion: A study on the robustness and generalization. IEEE Transactions on Geoscience and Remote sensing, 58(10):6900–6913, 2020.
- [42] Chengyuan Deng, Yanan Feng, Shihang Feng, Peng Jin, Xitong Zhang, Qili Zeng, and Youzuo Lin. OpenFWI: Benchmark seismic datasets for machine learning-based full waveform inversion. arXiv preprint arXiv:2111.02926, 2021.
- [43] Lukas Mosser, Olivier Dubrule, and Martin J Blunt. Stochastic seismic waveform inversion using generative adversarial networks as a geological prior. Mathematical Geosciences, 52(1):53–79, 2020.
- [44] Alan Richardson. Generative adversarial networks for model order reduction in seismic full-waveform inversion. arXiv preprint arXiv:1806.00828, 2018.
- [45] Weiqiang Zhu, Kailai Xu, Eric Darve, Biondo Biondi, and Gregory Beroza. Integrating deep neural networks with full-waveform inversion: Reparametrization, regularization, and uncertainty quantification. Geophysics, 87:1JF – WA102, 2021.
- [46] Yulang Wu and George A McMechan. Parametric convolutional neural network-domain full-waveform inversion. Geophysics, 84(6):R881–R896, 2019.
- [47] Qinglong He and Yanfei Wang. Reparameterized full-waveform inversion using deep neural networks. Geophysics, 86(1):V1–V13, 2021.
- [48] Yuxiao Ren, Xinji Xu, Senlin Yang, Lichao Nie, and Yangkang Chen. A physics-based neural-network way to perform seismic full waveform inversion. IEEE Access, 8:112266 – 112277, 2020.
- [49] Wenyi Hu, Yuchen Jin, Xuqing Wu, and Jiefu Chen. Progressive transfer learning for low-frequency data prediction in full-waveform inversion. Geophysics, 86:1JA – X4, 2021.
- [50] Yuchen Jin, Wenyi Hu, Shirui Wang, Yuan Zi, Xuqing Wu, and Jiefu Chen. Efficient progressive transfer learning for full-waveform inversion with extrapolated low-frequency reflection seismic data. IEEE Transactions on Geoscience and Remote Sensing, 80, 2021.
- [51] Hongyu Sun and Laurent Demanet. Extrapolated full-waveform inversion with deep learning. Geophysics, 85:1MJ – Z13, 2020.
- [52] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. arXiv preprint arXiv:1711.10925, 2020.

- [53] Wei Zhang and Jinghuai Gao. Deep-learning full-waveform inversion using seismic migration images. IEEE Transactions on Geoscience and Remote Sensing, 60, 2021.
- [54] Wei Zhang, Jinghuai Gao, Xiudi Jiang, and Wenbo Sun. Consistent least-squares reverse time migration using convolutional neural networks. IEEE Transactions on Geoscience and Remote Sensing, 60, 2021.
- [55] Wei Zhang, Jinghuai Gao, Zhaoqi Gao, and Hongling Chen. Adjoint-driven deep-learning seismic full-waveform inversion. IEEE Transactions on Geoscience and Remote Sensing, 59:8913–8932, 2020.
- [56] Yulang Wu, George McMechan, and Yanfei Wang. Adaptive feedback convolutional-neural-network-based high-resolution reflection-waveform inversion. Journal of Geophysical Research: Solid Earth, 127:e2022JB024138, 2022.
- [57] Yuqing Chen and Erdinc Saygin. Seismic inversion by hybrid machine learning. Journal of Geophysical Research: Solid Earth, 126:e2020JB021589, 2021.
- [58] Bingbing Sun and Tariq Alkhalifah. ML-misfit: Learning a robust misfit function for full-waveform inversion using machine learning. arXiv preprint arXiv:2002.03163, 2020.
- [59] Zhendong Zhang and Tariq Alkhalifah. Regularized elastic full-waveform inversion using deep learning. Geophysics, 84:R741 – R751, 2019.
- [60] United States Department of Energy. Energy Data eXchange (EDX). <https://edx.netl.doe.gov/>.
- [61] Oliver Hennigh, Susheela Narasimhan, Mohammad Amin Nabian, Akshay Subramaniam, Kaustubh Tangsali, Max Rietmann, Wonmin del Aguila Ferrandis, Joseand Byeon, Zhiwei Fang, and Sanjay Choudhry. NVIDIA SimNet: an ai-accelerated multi-physics simulation framework. arXiv preprint arXiv:2012.07938, 2020.
- [62] Lu Lu, Xuhui Meng, Zhiping Mao, and George Em Karniadakis. DeepXDE: a deep learning library for solving differential equations. SIAM Review, 63:208–228, 2021.
- [63] Alexander Koryagin, Roman Khudorozkov, and Sergey Tsimfer. PyDEns: a python framework for solving differential equations with neural networks. arXiv preprint arXiv:1909.11544, 2019.
- [64] Chengyuan Deng and Youzuo Lin. On the robustness and generalization of deep learning driven full waveform inversion. arXiv preprint arXiv:2111.14220, 2021.
- [65] Stephen James, Paul Wohlhart, Mrinal Kalakrishnan, Dmitry Kalashnikov, Alex Irpan, Julian Ibarz, Sergey Levine, Raia Hadsell, and Konstantinos Bousmalis. Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks. In IEEE International Conference on Computer Vision, 2019.
- [66] Yuntian Chen, Dou Huang, Dongxiao Zhang, Junsheng Zeng, Nanzhe Wang, Haoran Zhang, and Jinyue Yan. Theory-guided hard constraint projection (HCP): A knowledge-based data-driven scientific machine learning method. Journal of Computational Physics, 445, 2021.
- [67] Evan Schankee Um, David Alumbaugh, Youzuo Lin, and Shihang Feng. Real-time deep-learning inversion of seismic full waveform data for CO<sub>2</sub> saturation and uncertainty in geological carbon storage monitoring. Geophysical Prospecting, 2022 (Accepted).
- [68] Yinan Feng, Yinpeng Chen, Shihang Feng, Peng Jin, Zicheng Liu, and Youzuo Lin. An intriguing property of geophysics inversion. In The Thirty-ninth International Conference on Machine Learning (ICML), 2022.
- [69] Qili Zeng, Shihang Feng, Brendt Wohlberg, and Youzuo Lin. InversionNet3D: Efficient and scalable learning for 3-D full-waveform inversion. IEEE Transactions on Geoscience and Remote Sensing, 60:1–16, 2022.
- [70] Evan Schankee Um, David Alumbaugh, Michael Commer, Shihang Feng, Erika Gasperikova, Yaoguo Li, Youzuo Lin, and Savini Samarasinghe. Deep-learning multiphysics network for imaging CO<sub>2</sub> saturation and estimating uncertainty in geological carbon storage. Geophysical Prospecting, 2022 (Accepted).

- [71] Shihang Feng, Peng Jin, Yinpeng Chen, Xitong Zhang, Zicheng Liu, and Youzuo Lin. Exploring multi-physics with extremely weak supervision. arXiv preprint arXiv:2202.01770, 2022.
- [72] Vladimir Kazei, Oleg Ovcharenko, Pavel Plotnitskii, Daniel Peter, Xiangliang Zhang, and Tariq Alkhalifah. Mapping full seismic waveforms to vertical velocity profiles by deep learning. Geophysics, 86:1SO–Z1, 2021.