Physics-Consistent Data-driven Seismic Inversion with Adaptive Data Augmentation

Renán Rojas-Gómez
Los Alamos National Laboratory
Los Alamos, NM 87544

Jihyun Yang
Los Alamos National Laboratory
Los Alamos, NM 87544

Youzuo Lin (corresponding author)
Los Alamos National Laboratory
Los Alamos, NM 87544
ylin@lanl.gov

James Theiler
Los Alamos National Laboratory
Los Alamos, NM 87544

Brendt Wohlberg
Los Alamos National Laboratory
Los Alamos, NM 87544

Abstract

Solving the seismic full-waveform inversion (FWI) problem can be challenging due to its ill-posedness and high computational cost. We develop a new hybrid computational approach to solve FWI that combines physics-based models with data-driven methodologies. In particular, we develop a data augmentation strategy that can not only improve the representativity of the training set, but also incorporate important governing physics into the training process and therefore improve the inversion accuracy. We demonstrate our method with an example of monitoring subsurface carbon sequestration leakage. Our method yields higher accuracy and greater generalization ability than purely physics-based and purely data-driven approaches.

1 Introduction

Seismic full-waveform inversion (FWI) attempts to reconstruct an image of the subsurface geology from measurements of natural or artificially produced seismic waves that have travelled through the subsurface. The seismic FWI problem is challenging due to the non-linearity of the forward model and its under-determined nature. Conventional computational methods for solving FWI are based on optimization techniques and generic regularization [11]. For simplicity of description, we call these approaches "physics-based FWI methods" to distinguish them from data-driven methods, and from our proposed hybrid approach. The major advantage of these physics-based methods is their robustness to out-of-distribution data, while disadvantages include computational expense and the need for explicitly-formulated generic regularization.

Here, we describe a data-driven approach for seismic FWI that incorporates the physics model into the learning procedure. Specifically, this physics-consistent data-driven full waveform inversion consists of a carefully designed encoder-decoder-structured neural network and an adaptive data

Third Workshop on Machine Learning and the Physical Sciences (NeurIPS 2020), Vancouver, Canada.

augmentation technique. This augmentation employs the forward model to produce new training data that are more representative of the solution we seek. To validate its performance, we apply our inversion method to detect carbon sequestration leakage using synthetic seismic data sets generated using a subsurface model for a potential CO_2 storage site at Kimberlina, California [2].

Related Work. A new class of algorithm has recently emerged, based on machine learning applied to large datasets that are produced from many runs of the forward physics model. In direct end-to-end learning [1, 3, 5, 12, 13, 14], a large number of velocity maps and corresponding seismic waveforms (usually constructed through extensive simulation) are used as training data in learning the mapping from seismic waveform to velocity map. In low-wave number learning [6, 10], this type of learning approach is used to predict an initial velocity map with low-frequency component, which is then used as the initial guess for traditional physics-based optimization.

2 Background

Governing Physics: the Forward Model. Mathematically, the forward model can be expressed in terms of the seismic elastic-wave partial differential equation [11]:

$$\rho(r)\frac{\partial^2 u(r,t)}{\partial t^2} = (\lambda(r)+\mu(r))\nabla(\nabla\cdot u(r,t)) + \mu(r)\nabla^2 u(r,t) + s(r,\,t), \tag{1}$$

where $\rho(r)$ is the density at spatial location r, $\lambda(r)$ and $\mu(r)$ are the Lamé parameters, $s(r,\,t)$ is the source term, $u(r,t)$ is the displacement wavefield, t represents time, and $\nabla\cdot$ is the divergence operator. When fluid such as supercritical CO_2 leaks into the subsurface formation, the geophysical parameters of P-wave and S-wave velocities will be changed correspondingly. Instead of inverting for $\rho(r)$, $\lambda(r)$ and $\mu(r)$, it is customary to invert for a velocity map $m(r)$ that depends on $\rho(r)$, $\lambda(r)$ and $\mu(r)$. We discretize r so that the velocity map we seek is $m \in \mathbb{R}^{M\times N}$, where M and N are its vertical and lateral dimensions, respectively. Here m refers to either P-wave or S-wave velocity. Similarly, we denote a seismic data observation $d_{\text{obs}} \in \mathbb{R}^{T\times S\times R}$, where T corresponds to the number of samples in the temporal domain, S to the number of sources and R to the number of receivers used in the data acquisition process. The seismic data can be expressed in terms of a highly nonlinear forward mapping:

$$d_{\text{obs}} = f(m), \tag{2}$$

where the forward operator f represents the wave propagation as provided in Eq. (1).

Physics-Based Inversion. Various explicit regularization techniques have been developed to stabilize the computation of seismic inversion, including ℓ_1-norm [9] and ℓ_2-norm [4] methods. Given the forward model in Eq. (2), the regularized seismic FWI can be posed as

$$m = \underset{m}{\operatorname{argmin}} \left\{ \left\| d - f(m) \right\|_2^2 + \lambda\, R(m) \right\}, \tag{3}$$

where d represents a recorded waveform dataset, $f(m)$ is the corresponding forward modeling result, $\left\| d - f(m) \right\|_2^2$ is the data misfit term, λ is a regularization parameter and $R(m)$ is the regularization term. The formulation implicitly assumes the Gaussian noise being imposed on the data.

Data-Driven Inversion. A data-driven FWI structure based on an *encoder-decoder* architecture [12], denoted $\mathcal{G}$ and characterized by hyperparameters $\boldsymbol{\theta}$, is proposed to approximate the inverse mapping f^{-1} and obtain accurate velocity map predictions $\hat{m}(\boldsymbol{\theta}) \triangleq \mathcal{G}(\boldsymbol{\theta}, d_{\text{obs}})$. Optimal parameters $\boldsymbol{\theta}^*$ are obtained by adapting the architecture to a representative training set with L samples $\{d_{\text{obs},\ell}, m_\ell^*\}, \ell \in \{0, L-1\}$. We choose the mean-absolute error (MAE) as our optimality criterion:

$$\boldsymbol{\theta}^* = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \frac{1}{L}\sum_{\ell=0}^{L-1} \|m_\ell^* - \hat{m}_\ell(\boldsymbol{\theta})\|_1. \tag{4}$$

For a more detailed discussion of loss function, please refer to our earlier work [12].

3 Physics-Consistent Data-Driven Full-waveform Inversion

Data Description. We apply our method to detect CO_2 leakage in the subsurface. Because direct subsurface geological measurements are scarce to non-existent, we use the simulated Kimberlina

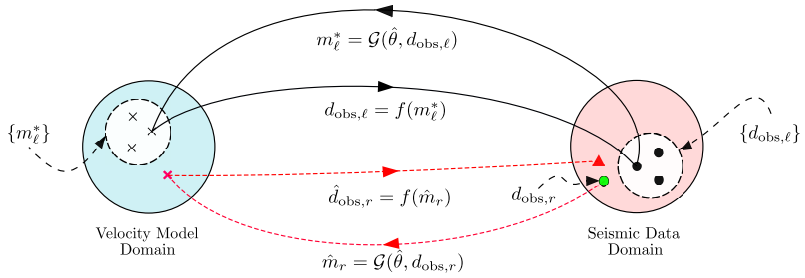

Figure 1: Adaptive Data Augmentation: Approximate solver $\mathcal{G}(\hat{\theta}, d_{\text{obs}})$ is fully-trained over labeled set $\{m_\ell^*, d_{\text{obs},\ell}\}$, and applied to unlabeled seismic data $d_{\text{obs},r}$ to generate new velocity maps $\hat{m}_r$. Physically-coherent seismic data $\hat{d}_{\text{obs},r}$ is then generated using the forward model f, producing a new labeled set $\{\hat{m}_r, \hat{d}_{\text{obs},r}\}$ which is added to the original training set.

dataset from Lawrence Livermore National Laboratory (and that we refer to here as `CO2leak`). The aim of the Kimberlina dataset is to understand and assess the effectiveness of various geophysical monitoring techniques in detecting CO_2 shallow leakage in the wellbore [7]. The `CO2leak` dataset contains 991 CO_2 leakage scenarios, each simulated over a duration of 200 years, with 20 leakage maps provided (ie, at every ten years) for each scenario. We obtain synthetic seismograms from elastic forward modeling on `CO2leak` velocity maps. First, one-second traces with a time interval of 0.5ms using 7 sources and 114 receivers are generated. We then down-sample each trace by a factor of 2, resulting in a temporal dimension of 1000 time steps. The sources and receivers are evenly distributed along the top of the model, with depths of 5m and 20m, respectively. The source interval is 125m, and the receiver interval is 15m. We use a Ricker wavelet with a central frequency of 25Hz as the source to generate simulated seismic waves due to its empirical success in processing field data [4]. The synthetic data is the staggered grid solution of the elastic wave equation using a finite-difference scheme with a perfectly matched layered absorbing boundary condition.

Data Augmentation: Incorporation of Physics Knowledge. The `CO2leak` dataset includes $19,600$ velocity maps of 141×341 grid points describing CO_2 and brine leakage plumes evolving with time. It is desirable to detect plumes of leaking CO_2 while they are still small. This is particularly challenging when the available training data is dominated by large plumes. Thus, `CO2leak` presents the opportunity to evaluate the *generalization* of data-driven method regarding different plume sizes.

Along with the data pairs included in the dataset, the ground-truth CO_2 and brine mass information for each sample, is provided. Based on this, the full dataset is split into four parts, according to their CO_2 leak mass plus brine leak mass: *tiny* plumes (from 3.53×10^2 to 9.10×10^6Kg), *small* plumes (from 9.10×10^6 to 2.67×10^7Kg), *medium* plumes (from 2.67×10^7 to 8.05×10^7Kg), and *large* plumes (from 8.05×10^7 to 1.62×10^9 Kg). These cover 20%, 20%, 20%, and 40% of the data samples, respectively. While conventional data augmentation techniques (such as rotation, flip, scale etc.) have proved to be effective for computer vision applications, it is not clear that they have a useful role to play in our application because those invariances do not apply for this problem. Our adaptive data augmentation scheme provides additional training data that is not only physically meaningful but also more closely related to the target unlabelled data that we are trying to invert. We summarize our augmentation method with the following four steps (illustrated in Fig. 1):

i. Estimate approximate solver $\mathcal{G}(\hat{\boldsymbol{\theta}}, d_{\text{obs}})$;

ii. Generate approximate velocity maps from unlabeled data $\hat{m}_r = \mathcal{G}(\hat{\boldsymbol{\theta}}, d_{\text{obs},r})$;

iii. Create seismic data using forward model $\hat{d}_{\text{obs},r} = f(\hat{m}_r)$;

iv. Add new pairs $(\hat{d}_{\text{obs},r}, \hat{m}_r)$ to the original training set.

The augmented dataset plays a key role in model accuracy because it will not only carry useful physics information, but also provides examples of velocity maps that are consistent with the target geology feature of interests. Furthermore, the full augmentation process can be applied in an iterative fashion by re-training the approximate solver $\mathcal{G}(\hat{\boldsymbol{\theta}}, d_{\text{obs}})$ based on the extended training set in order to generate new approximate velocity maps $\hat{m}_r$. This approach allows further refinement of the mapping between velocity and seismic subdomains.

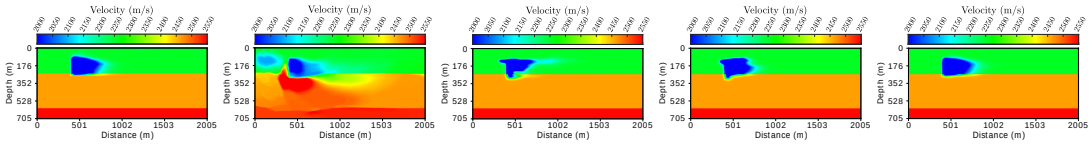

(a) (b) (c) (d) (e)

Figure 3: (a) Ground truth. Inversions obtained using (b) physics-based FWI, data-driven model trained over (c) the *large* subset, (d) *large* and *augmented medium* subsets (Augmented once), and (e) *large* and *augmented medium* subsets (Augmented twice).

4 Numerical Experiments

Experimental Setup. The training process is performed using a fixed stopping criterion of 250 epochs, random weight initialization, an initial learning rate of 10^{-3} and its subsequent adaptive optimization via ADAM [8]. The training is performed using fixed batch sizes of 50, where both features and labels are normalized prior to their use in the training and inference processes. We created 392 batches in total. In consequence, reported loss values are computed based on normalized training and testing pairs. All training and testing routines are implemented using *PyTorch* code and executed on four *NVIDIA GeForce GTX 1080* GPUs.

Generalization Performance with Data Augmentation. In our test, let the *medium unlabeled* subset correspond exclusively to the *medium* seismic data (the corresponding velocity maps are not taken into account). Then, consider the following two experiments with different training set while applying to the same *medium* subset for test:

Test 1: Train on the *large* subset.

Test 2: Train on the *large* and augmented *medium unlabeled* subset.

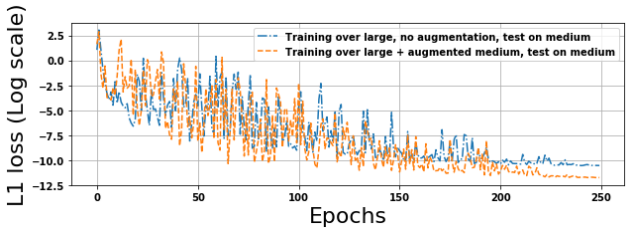

Figure 2: Reconstruction results for our data Augmentation approach applied to unlabeled data: Testing MAE loss value for each training epoch.

The augmentation process follows the procedure described in Fig. 1. Only the augmented *medium* dataset is added to the initial *large* dataset. Figure 2 shows reconstruction results along 250 epochs. Given that some of the plumes are much smaller than the full velocity map dimensions, we show the reconstruction accuracy along epochs $\varepsilon(\boldsymbol{\theta}_i)$ as the MAE in logarithmic scale and normalized with respect to the velocity map dimensions to clearly depict our method's effect in the predictions $\varepsilon(\boldsymbol{\theta}_i) = 10\log_{10}\left\{\frac{1}{\hat{L}MN}\sum_{\ell=0}^{\hat{L}-1}\|m_\ell^* - \hat{m}_\ell(\boldsymbol{\theta}_i)\|_1\right\}$, where $\boldsymbol{\theta}_i$ corresponds to the set of hyperparameters at the i^{th} epoch, and $\hat{L}$ corresponds to the size of the validation set. The network trained over *large* plumes attains a reconstruction accuracy of $\varepsilon \approx -10.5$. On the other hand, the network trained over *large* and augmented *medium unlabeled* plumes obtain a better reconstruction accuracy of approximately -11.7. These results strengthen our observation regarding the information encapsulated in the samples generated by the augmentation process: by including the forward modeling operation in the process, physically-consistent data pairs are generated, allowing a better domain adaptation. The use of unlabeled data shows how our proposed method is not limited to labeled data, which can be difficult to obtain in real applications. We provide visualization of the inverted velocity maps using both physics-based FWI and data-driven approaches and demonstrate its accuracy in Fig. 3. Our network trained over the *large* dataset obtains a reasonably accurate estimate of the plume. Also, the estimate obtained by our augmentation further refines the reconstruction shape and the plume location. This example shows how iterative augmentation can improve the network output, which reflects the potential of incorporating the physics-based forward model into the learning pipeline. An

additional test and analysis of our inversion model robustness with respect to variant levels of noisy seismic data is provided in our full paper [5].

5 Conclusion

We develop a physics-consistent data-driven seismic inversion method. We design a novel data augmentation strategy that incorporates critical physics information and improves the representability of the training set. We validate its performance to detect CO_2 leakage. Compared with purely physics-based and purely data-driven inversion methods, our physics-consistent data-driven inversion yields higher accuracy and better generalization. More detail of our work can be found in the full paper [5].
Acknowledgements. This work was supported by the Center for Space and Earth Science at Los Alamos National Laboratory (LANL) and by the Laboratory Directed Research and Development program of LANL under project numbers of 20210542MFR and 20200061DR.

Broader Impact

Inverse problems, the inference of unknown or unobservable properties from measurements, have a diverse range of scientific applications including materials science, medicine, and geoscience. The solutions to scientific inverse problems play an important role in characterizing the complexity of physics system and guiding future data acquisition. However, inverse problems are notoriously formidable: ill-posed and computationally expensive. Our developed approach provides a mean in exploiting machine learning in a way enhances, rather than replaces, traditional physics-based inversion methods. The idea does not only limit itself only to our geophysics problems, but instead, it can be potentially applied to much broader inverse problems.

References

[1] Araya-Polo, M., J. Jennings, A. Adler, and T. Dahlke, 2018, Deep-learning tomography: The Leading Edge, **37**, 58–66.

[2] Buscheck, T., K. Mansoor, X. Yang, H. Wainwright, and S. Carroll, 2019, Downhole pressure and chemical monitoring for CO_2 and brine leak detection in aquifers above a CO_2 storage reservoir: Int. J. Greenhouse Gas Control, **91**.

[3] Farris, S., M. Araya-Polo, J. Jennings, B. Clapp, and B. Biondi, 2018, Tomography: a deep learning vs full-waveform inversion comparison, *in* First EAGE Workshop on High Performance Computing for Upstream in Latin America, European Association of Geoscientists & Engineers.

[4] Fichtner, A., 2010, Full seismic waveform modelling and inversion: Springer Science & Business Media.

[5] Gomez, R., J. Yang, Y. Lin, J. Theiler, and B. Wohlberg, 2020, Physics-consistent data-driven waveform inversion with adaptive data augmentation: arXiv preprint (also accepted in IEEE Geoscience and Remote Sensing Letters).

[6] Hu, W., Y. Jin, X. Wu, and J. Chen, 2019, Progressive transfer learning for low frequency data prediction in full waveform inversion: arXiv preprint arXiv:1912.09944.

[7] Jordan, P., and J. Wagoner, 2017, Characterizing construction of existing wells to a CO_2 storage target: The Kimberlina site, California: Technical report, U.S. Department of Energy - Office of Fossil Energy.

[8] Kingma, D. P., and J. Ba, 2014, Adam: A method for stochastic optimization: arXiv preprint arXiv:1412.6980.

[9] Lin, Y., and L. Huang, 2015, Acoustic- and elastic-waveform inversion using a modified total-variation regularization scheme: Geophysical Journal International, **200**, 489–502.

[10] Ovcharenko, O., V. Kazei, M. Kalita, D. Peter, and T. Alkhalifah, 2019, Deep learning for low-frequency extrapolation from multioffset seismic data: Geophysics, **84**, no. 6, 58–66.

[11] Virieux, J., A. Asnaashari, R. Brossier, L. Métivier, A. Ribodetti, and W. Zhou, 2014, Chapter 6: An introduction to full waveform inversion, *in* Encyclopedia of Exploration Geophysics: SEG.

[12] Wu, Y., and Y. Lin, 2019, InversionNet: An efficient and accurate data-driven full waveform inversion: IEEE Transactions on Computational Imaging, **6**, 419–433.

[13] Yang, F., and J. Ma, 2019, Deep-learning inversion: A next-generation seismic velocity model building method: Geophysics, **84**, no. 4, R583–R599.

[14] Zhang, Z., and Y. Lin, 2020, Data-driven seismic waveform inversion: A study on the robustness and generalization: IEEE Transactions on Geoscience and Remote Sensing, **58**, no. 10, 6900–6913.