Tutorial on Novel Toolkits toward AI for Science on Resource-Constrained Computing Systems

Yi Sheng[†], Junhuan Yang[†], Hanchen Wang[§‡], Yinan Feng[‡], Yinpeng Chen[*],
Xiaolong Guo[¶], Youzuo Lin[‡], Weiwen Jiang[†], Lei Yang[†]
[†]George Mason University [§]Los Alamos National Laboratory [‡]University of North Carolina at Chapel Hill
[*]Google [¶]Kansas State University
{ysheng2, lyang29}@gmu.edu

I. ABSTRACT

Full Waveform Inversion (FWI) is a technique used to visualize and analyze wave propagation through a medium in order to infer its physical properties. This method relies on computational models and algorithms to simulate and interpret the behavior of waves—such as sound, electromagnetic, or seismic waves—as they travel through different materials. By analyzing how these waves are reflected, refracted, or absorbed by the medium, FWI can provide detailed information about the medium's internal structure, composition, and physical properties, such as density, elasticity, or internal defects. The traditional process typically involves: 1) Wave Simulation: Using physics-based models to simulate how waves propagate through a medium. This may involve solving complex differential equations that describe wave behavior in different contexts. 2) Data Acquisition: Collecting data on wave interactions with the medium using sensors or other measurement devices. This could include data on wave speed, direction, amplitude, and phase changes. 3) Image Reconstruction: Applying computational techniques, such as inverse problems or tomographic reconstruction, to create images or maps of the medium based on the acquired wave data. 4) Analysis: Interpreting the reconstructed images to deduce the physical properties of the medium. This can involve identifying features like boundaries, interfaces, or anomalies within the medium.

A. *Application and dataset*

FWI has applications in various fields, including medical imaging (like ultrasound), geophysics (like seismic imaging for oil exploration), and materials science (for non-destructive testing and evaluation of materials). However, physics-based models always meet challenges. Firstly, the expensive inversion operator. An "inversion" operator refers to determining a medium's internal properties from measurements of how waves propagate through it. This process often involves solving an inverse problem, where the goal is to infer the medium's properties (such as density or elasticity) from observed data (like wave travel times or amplitudes). The high computational cost will associated with the inversion process. Secondly, the challenges in accurately resolving the fine details of the medium's internal structure. This limitation can arise due to several factors, including poor data coverage, severe ill-posedness of the inverse problem, and the use of simplistic regularization techniques. Facing the challenges of the physics-based models, machine learning (ML) approaches have been applied to FWI, which significantly compensates for the high computational resource demands and lengthy processing times.

In this tutorial, we will first introduce the OpenFWI dataset [1], which is a collection of large-scale, multi-structural benchmark datasets for a machine learning-driven seismic FWI. We release twelve datasets synthesized from different priors, including one 3D dataset. We also provide baseline experimental results with four deep learning methods: InversionNet [2], VelocityGAN [3], UPFWI [4] and InversionNet3D [5].

On top of the fundamental datasets and models, three recent toolkits will be introduced with hands-on experience, including *EdGeo* [6] for geophysics data generation on edge devices, *APS-USCT* [7]–[9] for reconstructing the human body using very sparse ultrasonic waves, and *QuGeo* [10] for geophysics processing on resource-constrained noisy intermediate-scale quantum (NISQ) devices.

B. *Toolkit 1: EdGeo*

Machine learning (ML) comes with inherent challenges. Unlike physics-driven methods, which can be universally applied across diverse locations and conditions, ML often performs poorly on unprivileged data. This issue arises because of the diversity and complexity of subsurface structures in different locations, as well as dynamic changes in underground conditions (e.g., petroleum leaks), making unprivileged data a common occurrence in geoscience. However, most previous data-driven approaches have designed ML models without addressing this issue. Consequently, pre-trained ML models often fail to perform effectively in geoscience applications, necessitating a process of localization to handle unprivileged data, which severely limits the effectiveness of ML. This issue is further exacerbated during model pruning, a crucial step in geoscience due to environmental complexities.

To tackle these challenges, the EdGeo toolkit [6] employs a diffusion-based model guided by physics principles to generate high-fidelity velocity maps. It uses the acoustic wave equation to produce the corresponding seismic waveform data, which facilitates the fine-tuning of pruned ML models. The proposed EdeGo has 2 stages: offline and online. The offline phase utilizes seismic data and the corresponding velocity maps to pre-train a pruned InversionNet model. The velocity distribution across different layers is determined based on unprivileged data or expert experience. The online phase comprises six

2832-6474/24/$31.00 ©2024 IEEE
DOI 10.1109/CODES-ISSS60120.2024.00006

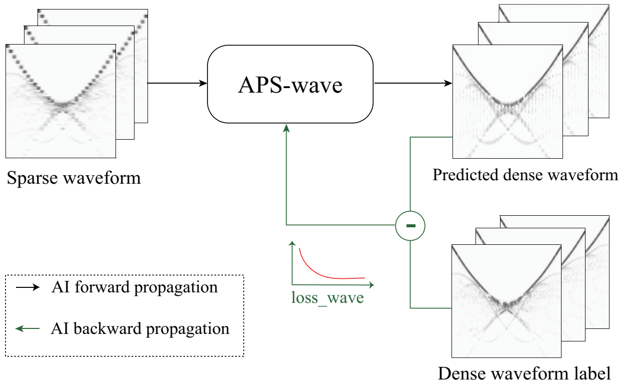

Fig. 1. Overview of the proposed APS-USCT

modules and methods designed to generate velocity maps and seismic data under specific conditions. Our approach is specifically designed for real-world applications. It focuses on real-time processing and adherence to resource constraints, which ensures effective localization of the ML model. Our results show significant improvements in SSIM scores and reductions in MAE and MSE across various pruning ratios.

C. Toolkit 2: APS-USCT

Although efficient, machine learning performance (i.e., image quality) depends heavily on highly dense waveforms, which necessitate expensive equipment with numerous transducers (i.e., sources and receivers).

To address this challenge, we propose exploring the feasibility of achieving high-quality reconstructed images by enhancing the available sparse data through an AI-physics synergy framework. The framework begins by upscaling sparse waveforms using an AI approach referred to as APS-wave [7] (illustrated in Fig 1) to generate dense waveforms. This process is enabled by constructing a training dataset of dense waveforms through APS-physics. Next, the generated dense waveform is processed by the second AI component, called APS-FWI, which employs InversionNet as its backbone architecture, enhanced with SE-Blocks and source encoding. The SE-Blocks improve the capture of fine details in the reconstruction of the Speed of Sound (SOS) map, while the source encoding increases the model's learning efficiency. In the framework, the AI module (APS-wave) and the physics module (APS-physics) collaborate to transform sparse measurements into dense waveforms, thereby increasing sample density before reconstruction, which enhances data density while preserving waveform integrity.

We tested APS-USCT on a breast reconstruction dataset, where it significantly outperformed state-of-the-art techniques. Compared to the leading approach using dense input waveforms, we can achieve a 2.5× reduction in hardware costs (i.e., fewer transducers) with only a minor SSIM degradation.

D. Toolkit 3: QuGeo

Realizing the potential of quantum computing hinges on identifying "killer applications". We introduce QuGeo [10], an innovative quantum framework to address FWI challenges.

In this tutorial, we will first introduce a physics-informed dataset using a governing wave equation, upon which we developed a classical machine learning-based data converter. This converter efficiently scales the data according to quantum resource constraints. Secondly, we will present an application-specific Variational Quantum Circuit (VQC), called QuGeoVQC, which integrates domain-specific knowledge to optimize its design. QuGeoVQC focuses on the design of a data encoder and VQC computing structure to extract spatial and temporal features. It also leverages the unique characteristics of the FWI problem to simplify and optimize the encoder, thereby enhancing performance. Additionally, we will introduce a novel data batching technique tailored for quantum computing. This technique allows quantum computers to process N batches of data in parallel with only an additional log_N qubits, effectively meeting the high computational demands of learning-based FWI and advancing QuGeo's potential in seismic inversion.

To evaluate the effectiveness, we conducted tests using OpenFWI. Our results demonstrate that incorporating physics knowledge into QuGeo not only achieves high prediction accuracy, as indicated by SSIM values, but also results in an efficient learning model that utilizes only 576 parameters.

ACKNOWLEDGEMENT

We gratefully acknowledge the support of the National Institutes of Health (NIH) (Award No. 1R01EB033387-01), National Science Foundation (NSF) 2311949, and NSF 2320957.

REFERENCES

[1] C. Deng, S. Feng, H. Wang, X. Zhang, P. Jin, Y. Feng, Q. Zeng, Y. Chen, and Y. Lin, "Openfwi: Large-scale multi-structural benchmark datasets for full waveform inversion," in *Advances in Neural Information Processing Systems*, vol. 35, Curran Associates, Inc., 2022.

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

[3] Z. Zhang, Y. Wu, Z. Zhou, and Y. Lin, "Velocitygan: Subsurface velocity image estimation using conditional adversarial networks," in *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 705–714, IEEE, 2019.

[4] P. Jin, X. Zhang, Y. Chen, S. X. Huang, Z. Liu, and Y. Lin, "Unsupervised learning of full-waveform inversion: Connecting cnn and partial differential equation in a loop," *arXiv preprint arXiv:2110.07584*, 2021.

[5] Q. Zeng, S. Feng, B. Wohlberg, and Y. Lin, "Inversionnet3d: Efficient and scalable learning for 3-d full-waveform inversion," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–16, 2021.

[6] J. Yang, H. Wang, Y. Sheng, Y. Lin, and L. Yang, "A physics-guided generative ai toolkit for geophysical monitoring," *In Proc. of ACM/IEEE Design Automation Conference (DAC)*, 2024.

[7] Y. Sheng, H. Wang, Y. Liu, J. Yang, W. Jiang, Y. Lin, and L. Yang, "Aps-usct: Ultrasound computed tomography on sparse data via aiphysic synergy," *In Proc. of International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2024.

[8] Y. Sheng, J. Yang, Y. Lin, W. Jiang, and L. Yang, "Toward fair ultrasound computing tomography: Challenges, solutions and outlook," in *Proceedings of the Great Lakes Symposium on VLSI 2024*, pp. 748–753, 2024.

[9] J. Yang, Y. Sheng, Y. Zhang, H. Wang, Y. Lin, and L. Yang, "Enhanced ai for science using diffusion-based generative ai-a case study on ultrasound computing tomography," in *Proceedings of the Great Lakes Symposium on VLSI 2024*, pp. 754–759, 2024.

[10] W. Jiang and Y. Lin, "Qugeo: An end-to-end quantum learning framework for geoscience–a case study on full-waveform inversion," *In Proc. of ACM/IEEE Design Automation Conference (DAC)*, 2024.