AI-Driven Void Detection Using SPECFEM3D Wave Propagation Simulations

SINA TAJMIRI, BOYOUNG KIM and CHANSEOK JEONG*

ABSTRACT

This study introduces a deep learning framework for detecting underground cavities using synthetic seismic displacement data generated by SPECFEM3D. Traditional inversion methods are computationally intensive and often limited in resolution, motivating the use of a 3D convolutional neural network (3DCNN) to learn spatial-temporal features from sparse surface sensor data. The model was trained to predict a continuous 3D probability matrix indicating cavity presence, which was later post-processed for binary classification and evaluation. High accuracy was achieved across regression and classification metrics. Comparing simulated wave responses from predicted cavities to the original sensor data, confirmed physical consistency. Additional blind tests demonstrated robustness under varying conditions. Overall, this approach offers a promising, efficient alternative for subsurface anomaly detection with strong predictive performance and physical reliability.

INTRODUCTION

Subsurface anomalies, such as voids, which can pose serious risks in geotechnical engineering contexts, often produce subtle and complex wave disturbances that are difficult to interpret manually. Accurate identification of subsurface anomalies allows for proactive decision-making in construction, engineering, and geotechnical engineering, reducing risks and costs. As such, reliable detection methods are essential for maintaining both public safety and operational efficiency in subsurface-related applications. SPECFEM3D software, based on the spectral element method (SEM), can play a critical role in accurate wave propagation simulation and high-resolution 3D subsurface seismic imaging [1–3]. Additionally, SPECFEM3D incorporates perfectly matched layers

Sina Tajmiri, PhD Student, School of Engineering and Technology & Earth and Ecosystem Science, Cen- tral Michigan University, Mount Pleasant, MI, USA. Email: tajmi1s@cmich.edu
Boyoung Kim, Postdoctoral Researcher, School of Engineering and Technology, Central Michigan Uni- versity, Mount Pleasant, MI 48859, USA. Email: kim5b@cmich.edu
*Corresponding Author: Chanseok Jeong, Associate Professor, School of Engineering and Technol- ogy & Earth and Ecosystem Science, Central Michigan University, Mount Pleasant, MI, USA. Email: jeong1c@cmich.edu

(PMLs) [4–10] at the domain boundaries to effectively absorb outgoing waves, enabling realistic simulation of wave propagation in a semi-infinite medium without artificial boundary reflections. Leveraging its capability to generate realistic seismic responses in 3D media, this study addresses the challenge of interpreting complex wave response patterns caused by subsurface anomalies. While CNNs have shown great promise in spatial pattern recognition [11], their application to surface wave measurement-based cavity detection almost remains underexplored. Previous studies on subsurface anomaly detection often face limitations such as simplified 2D models, and poor generalization to complex or unseen geometries [12]. Many rely on manual feature extraction and struggle to accurately localize anomalies in 3D space [13], most previous studies rely on real-world data, which often lack ground truth about anomaly size, shape, and location. To address the limitations of machine learning-based anomaly imaging, element-wise classification [14, 15] were studied. Namely, anomaly imaging problems were formulated as classification problem such that a trained neural network classify each element of a background mesh discretizing a domain as a void element or a non-void (intact) one. Building upon the robust performance of the element-wise classification method in a two-dimensional (2D) setting, this study leverage SPECFEM3D to generate realistic synthetic displacement-field data for given anomalies in a 3D domain and apply a 3DCNN capable of learning spatial features (e.g., cavity or non-cavity grids) within the domain directly from raw sensor inputs. Our approach demonstrates that CNNs can accurately interpret sensor data to localize and classify subsurface anomalies, offering a robust framework for geophysical interpretation and anomaly detection in mechanical applications, such as non-detractive test (NDT).

THE SPECFEM3D SIMULATION AND 3DCNN MODEL ARCHITECTURE

In this study, we utilized the SPECFEM3D solver to simulate elastic wave propagation in a 3D medium using SEM. The model features a rectangular cubic domain with homogeneous elastic properties and a surface array of 12×12 seismic sensors spaced at $60\ \mathrm{m}$ intervals. A point source at the center initiates wave propagation, with convolutional perfectly match layer (C-PML) to mimic a semi-infinite domain by absorbing boundary conditions and stress-free surface conditions to ensure realistic wave behavior and prevent artificial reflections [16]. The governing equations, based on works by Komatitsch and Tromp [17, 18], describe the displacement wavefield due to external forces and internal stresses in elastic media. The SEM approach uses a weak form of these equations, discretized into hexahedral elements with Gauss-Lobatto-Legendre (GLL) points for accurate integration and interpolation. This method enables efficient and precise modeling of wave interactions, including surface and body waves, across complex geometries.

As it depiced in Figure 1, we generated our training data by using SPECFEM3D Cartesian's built-in mesh generator to build a structured $960\ \mathrm{m} \times 960\ \mathrm{m} \times 420\ \mathrm{m}$ cubic model with $30\ \mathrm{m}$ hexahedral elements and three-layer C-PML absorbing boundaries. A 1.2 Hz Ricker wavelet source drove 1000 time steps ($\Delta t = 0.008\ \mathrm{s}$), which further down sampled to 500 time steps and all runs were performed on a single processor to avoid interpolation errors at subdomain interfaces, yielding high-fidelity synthetic waveforms for AI-based void detection. In data generation process, two Bash scripts fully

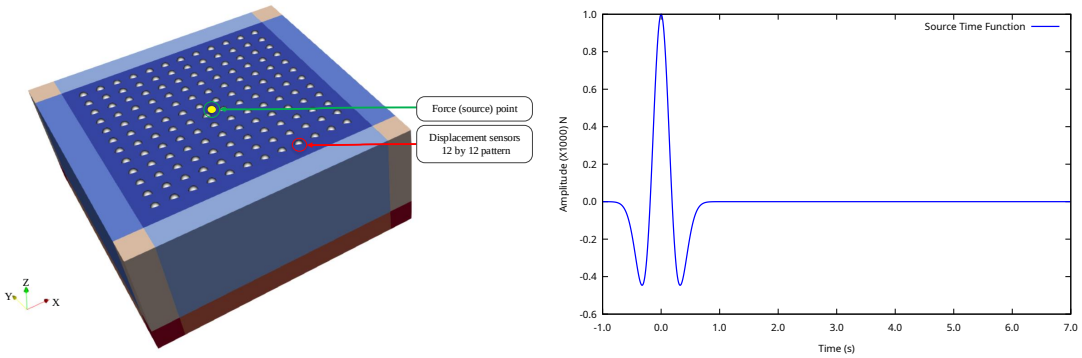

Figure 1. (left) SPECFEM3D model overview representing force source and the sensors' location and (right) force source time function plot.

automate the SPECFEM3D workflow by batch-updating cavity coordinates, running mesh-generation and solver steps with error checks and logging, and organizing each run's outputs, then extracting sensor displacement time-series, standardizing them via a z-score scaler ($z = \frac{x-\mu}{\sigma}$) and storing them in NumPy arrays. A companion randomizer generated 11,500 simulations in three groups, 4,500 single cavities, 5,000 sets of three non-overlapping cavities, and 2,000 sets with controlled overlaps to produce a diverse training dataset.

In SPECFEM3D, cavities are defined using start and end coordinates along the x, y, z axes, which are not ideal for training an AI model due to their sparse and discrete nature. To improve learning, we transformed these coordinates into a structured 3D binary matrix ($32 \times 32 \times 14$) marking cavity locations with ones and intact regions with zeros.

To tackle the imbalance between cavity and non-cavity points, we applied a proximity function, converting the binary labels into a smooth probability field based on distance to the cavity. This Gaussian-like function assigns higher values near cavities and gradually lowers them with distance, improving the CNN's ability to learn spatial patterns in the data. The CNN ingests spatio-temporal displacement tensors of shape (batch, 12, 12, 3, 500) and learns to predict a $32 \times 32 \times 14$ probability volume of cavity locations as shown in Figure 2. Data are split 80/10/10 for training, validation, and testing with a fixed random seed to ensure reproducibility. After an input layer and Gaussian noise, three convolutional blocks progressively extract temporal then spatial features: first, 64 filters of size $(1 \times 1 \times 5)$ followed by temporal max-pooling; next, 128 filters of $(1 \times 1 \times 3)$ with spatial downsampling to 6×6; and finally, 256 filters of $(3 \times 3 \times 1)$ with no pooling to capture localized anomalies. Each block uses batch normalization and increasing dropout (0.2→0.4). Extracted features are flattened and passed through four dense layers (512→256→128→64) with rectified linear unit (ReLU) activations and dropout, ending in a sigmoid output layer that yields the 3D probability field matrix. Training employs early stopping and learning-rate reduction on plateau, optimizing mean squared error (MSE) as loss function, and is evaluated via mean absolute error (MAE), and R^2 as:

$$\text{MSE} = \frac{1}{m}\sum_{j=1}^{m}\left(\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2\right)$$

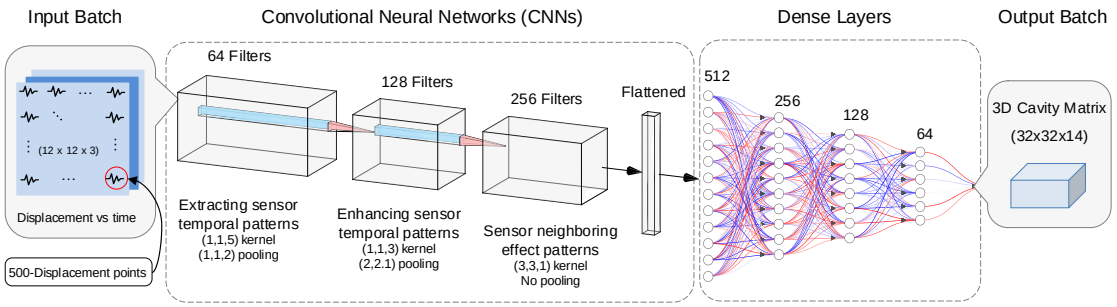

Figure 2. Schematic plot of the CNN process representing input, output, convolutional and pooling layers, and fully connected layers.

$$\text{MAE} = \frac{1}{m}\sum_{j=1}^{m}\left(\frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|\right), \; R^2 = 1 - \sum_{j=1}^{m}\left(\frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2}\right)$$

where n denotes the total number of components in a matrix array, m denotes number of batch data, y_i represents the true value from the ground truth 3D probability matrix at the i-th component index, while $\hat{y}_i$ is the corresponding predicted value from the model and $\bar{y}_i$ denotes the mean of the true value. The loss function (MSE) measures the average squared difference between these values, penalizing larger errors more heavily, which helps the model focus on correcting major discrepancies in cavity prediction. The MAE calculates the average of absolute differences, offering a more direct measure of prediction accuracy. The R^2 evaluates how well the model's predictions explain the variability in the actual data, indicating the overall goodness of fit. In the post-processing stage, the continuous 3D probability matrix predicted by the model is converted into a binary classification matrix to clearly identify cavity regions. This transformation is essential for practical applications that require a discrete yes/no classification rather than probabilistic estimates. A peak region finder function detects local maxima in the matrix, and only those exceeding a calibrated threshold—based on no-cavity simulations—are considered valid. These peaks are further expanded into surrounding regions using a prominence factor and Euclidean distance, forming smooth and realistic cavity shapes. This binary reconstruction allows for meaningful evaluation using following classification metrics, such as accuracy, precision, recall, and F_1-score, bridging the gap between probabilistic outputs and actionable cavity detection.

$$\textit{Accuracy} = \frac{TP + TN}{TP + FN + TN + FP} \times 100\ (\%), \; \textit{Precision} = \frac{TP}{TP + FP} \times 100\ (\%),$$

$$\textit{Recall} = \frac{TP}{TP + FN} \times 100\ (\%), \; \textit{F1-score} = \frac{2 \times \textit{Precision} \times \textit{Recall}}{\textit{Precision} + \textit{Recall}} \times 100\ (\%)$$

were the terms *TP*, *TN*, *FP*, and *FN* quantify the CNN's classification performance on the dataset. Specifically, *TP* measures correctly identified cavity elements, and *TN* measures correctly identified no-cavity elements. Conversely, *FP* counts no-cavity elements incorrectly flagged as cavity (errors of commission), while *FN* counts cavity elements missed by the classifier and labeled as no-cavity (errors of omission).

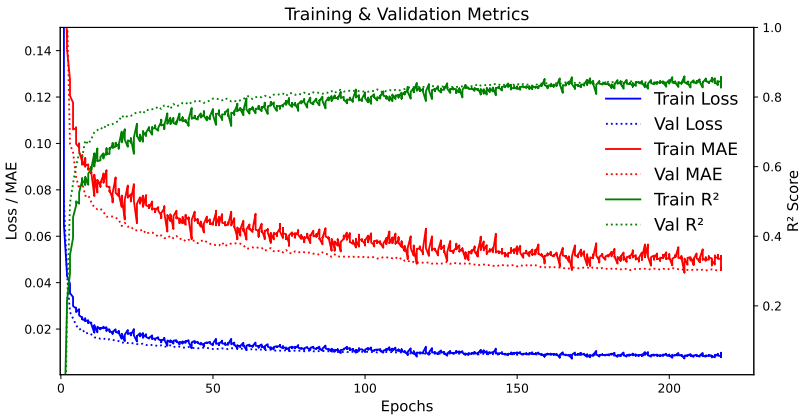

Figure 3. The values of loss, MAE and R^2 results over epochs during the training process.

TRAINING OF THE 3DCNN

The validation performance of the 3DCNN model was assessed by tracking training and validation loss, MAE, and R^2 over epochs. As ploted in Figure 3, the model exhibited rapid improvements in the first 50 epochs, with slower but stable gains thereafter. Training stopped at around 220 epochs based on early stopping criteria. Both training and validation MSE losses began below 0.03 and converged closely near 0.01, indicating strong generalization and no signs of overfitting or underfitting. Interestingly, MAE remained higher than MSE (around 0.05), with validation MAE slightly lower than training MAE possibly due to characteristics of the MSE calculations. The R^2 score quickly rose above 0.5 early on and stabilized near 0.85, with training and validation values closely aligned, suggesting that the model effectively captured data variance and maintained consistent predictive performance across both datasets.

NUMERICAL RESULTS

To further assess the generalization and robustness of the 3DCNN model, a blind test was conducted using complex, previously unseen cavity configurations not included in the training or validation datasets. These test cases featured irregular shapes, multiple anomalies, and varied depths to challenge the model's adaptability beyond the training distribution. Figure 4 illustrates a visual comparison between the true cavity geometries and the corresponding model predictions, highlighting the model's ability to approximate complex subsurface structures with high spatial accuracy. Despite the increased complexity, the predicted binary masks closely resemble the true cavity distributions, demonstrating strong spatial agreement in both shape and location. To quantitatively evaluate the model's classification performance, TABLE I presents key metrics including accuracy, precision, recall, and F_1-score for each test case. These metrics confirm the model's reliability, with consistently high accuracy scores across all blind tests, indicating effective detection of cavity regions with minimal false positives and false negatives. Together, the visual and numerical evaluations validate the model's capacity to generalize well up to 3 cavities to challenging scenarios, reinforcing its applicability in real-

TABLE I. Classification evaluation metrics of the prediction scenarios (Predictions 1 and 2 in Figure 4).

Var.	T_p	T_n	F_p	F_n	*Accuracy*	*Precision*	*Recall*	F_1-*score*
Prediction 1	30	14,336	14	6	99.86%	68.18%	83.33%	75.00%
Prediction 2	39	14,336	24	22	99.68%	61.90%	63.93%	62.90%

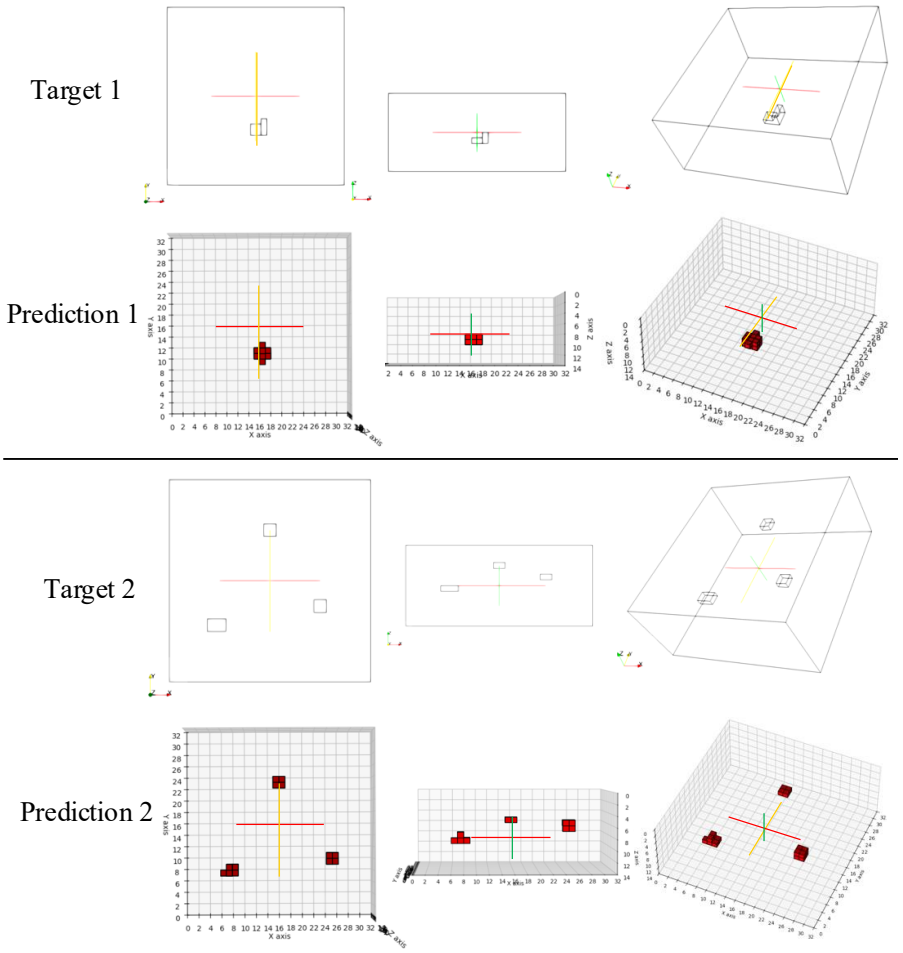

Figure 4. 3DCNN prediction performance visual comparison for two unseen target scenario (Predictions 1 and 2) with single and three cavities result.

world subsurface anomaly detection tasks. Figure 5 shows the displacement-field wave data at a subset of sensors computed for the predicted anomaly (Prediction 1) compared with their measured counterparts for original targets. The results showed strong agreement across all three spatial directions, with particularly close matches in the z-direction. Minor discrepancies in the x and y directions emerged after four seconds, likely due to slight geometric differences between the predicted and actual cavities. Overall, the similarity in waveforms confirms that the predicted cavity produces a physically consistent wavefield, validating the 3DCNN model's effectiveness and the practical relevance of its outputs.

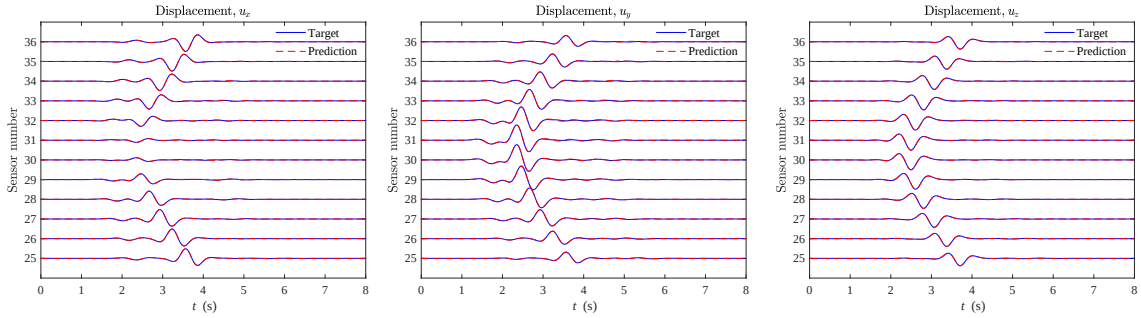

Figure 5. Displacement reconstruction test: comparison of measured wave data corresponding to targeted cavities and the reconstructed ones for the predicted cavities (Prediction 1).

CONCLUDING REMARKS

This work introduced an artificial neural network (ANN) framework for detecting subsurface voids leveraging SPECFEM3D elastic wave propagation simulation. The model demonstrated strong predictive accuracy (>99.5%) and generalization through both quantitative metrics and visual comparisons. Post-processing effectively converted probability outputs into actionable binary cavity maps, and blind tests confirmed the model's reliability on unseen cases. These results highlight the method's potential for practical geophysical and NDT applications.

REFERENCES

1. Chaljub, E., D. Komatitsch, J.-P. Vilotte, Y. Capdeville, B. Valette, and G. Festa. 2007. "Spectral-element analysis in seismology," *Advances in geophysics*, 48:365–419.
2. Komatitsch, D. and J. Tromp. 1999. "Introduction to the spectral element method for three-dimensional seismic wave propagation," *Geophysical journal international*, 139(3):806–822.
3. Peter, D., D. Komatitsch, Y. Luo, R. Martin, N. Le Goff, E. Casarotti, P. Le Loher, F. Magnoni, Q. Liu, C. Blitz, et al. 2011. "Forward and adjoint simulations of seismic wave propagation on fully unstructured hexahedral meshes," *Geophysical Journal International*, 186(2):721–739.
4. Berenger, J.-P. 1994. "A perfectly matched layer for the absorption of electromagnetic waves," *Journal of computational physics*, 114(2):185–200.
5. Francois, S., H. Goh, and L. F. Kallivokas. 2021. "Non-convolutional second-order complex-frequency-shifted perfectly matched layers for transient elastic wave propagation," *Computer Methods in Applied Mechanics and Engineering*, 377:113704.
6. Kang, J. W. and L. F. Kallivokas. 2010. "Mixed unsplit-field perfectly matched layers for transient simulations of scalar waves in heterogeneous domains," *Computational Geosciences*, 14(4):623–648.
7. Kucukcoban, S. and L. F. Kallivokas. 2011. "Mixed perfectly-matched-layers for direct transient analysis in 2D elastic heterogeneous media," *Computer methods in applied mechanics and engineering*, 200(1-4):57–76.
8. Fathi, A., B. Poursartip, and L. F. Kallivokas. 2015. "Time-domain hybrid formulations for wave simulations in three-dimensional PML-truncated heterogeneous media," *International Journal for Numerical Methods in Engineering*, 101(3):165–198.

9. Kim, B. and J. W. Kang. 2019. “A time-domain formulation of elastic waves in heterogeneous unbounded domains,” *Multiscale Science and Engineering*, 1:220–235.
10. Kim, B. and J. W. Kang. 2024. “Time domain modeling of elastic waves using a stress-based unsplit-field perfectly matched layer with enhanced numerical stability,” *Applied Mathematical Modelling*, 128:431–449.
11. Alzubaidi, L., J. Zhang, A. J. Humaidi, A. Al-Dujaili, Y. Duan, O. Al-Shamma, J. Santamaría, M. A. Fadhel, M. Al-Amidie, and L. Farhan. 2021. “Review of deep learning: concepts, CNN architectures, challenges, applications, future directions,” *Journal of big Data*, 8:1–74.
12. Zhang, G., Z. Wang, and Y. Chen. 2018. “Deep learning for seismic lithology prediction,” *Geophysical Journal International*, 215(2):1368–1387.
13. Bergen, K. J., P. A. Johnson, M. V. de Hoop, and G. C. Beroza. 2019. “Machine learning for data-driven discovery in solid Earth geoscience,” *Science*, 363(6433):eaau0323.
14. Pranto, F. M., S. Maharjan, and C. Jeong. 2023. “Level-Set and Learn: Convolutional Neural Network for Classification of Elements to Identify an Arbitrary Number of Voids in a 2D Solid Using Elastic Waves,” *Journal of Engineering Mechanics*, 149(6):04023035.
15. Kim, B., S. Maharjan, F. M. Pranto, B. Guidio, C. Schaal, and C. Jeong. 2024. “Convolutional neural network and level-set spectral element method for ultrasonic imaging of delamination cavities in an anisotropic composite structure,” *Ultrasonics*:107254.
16. Seriani, G. and E. Priolo. 1994. “Spectral element method for acoustic wave simulation in heterogeneous media,” *Finite elements in analysis and design*, 16(3-4):337–348.
17. Komatitsch, D. and J.-P. Vilotte. 1998. “The spectral element method: an efficient tool to simulate the seismic response of 2D and 3D geological structures,” *Bulletin of the seismological society of America*, 88(2):368–392.
18. Komatitsch, D. and J. Tromp. 2000. “The spectral element method for three-dimensional seismic wave propagation,” in *SEG International Exposition and Annual Meeting*, SEG, pp. SEG–2000.