
INVESTIGATING THE CAPABILITIES OF DEEP LEARNING FOR PROCESSING AND INTERPRETING ONE-SHOT MULTI-OFFSET GPR DATA: A NUMERICAL CASE STUDY FOR LUNAR AND MARTIAN ENVIRONMENTS

Iraklis Giannakis,

University of Aberdeen, School of Geosciences
Aberdeen, United Kingdom
iraklis.giannakis@abdn.ac.uk

Craig Warren

Northumbria University
Northumbria, United Kingdom
craig.warren@northumbria.ac.uk

Antonios Giannopoulos

The University of Edinburgh
Edinburgh, United Kingdom
a.giannopoulos@ed.ac.uk

Georgios Leontidis

Interdisciplinary Centre for Data and AI,
University of Aberdeen
Aberdeen, United Kingdom
georgios.leontidis@abdn.ac.uk

Yan Su

Chinese Academy of Sciences
Beijing, China
suyan@nao.cas.cn

Feng Zhou

China University of Geosciences (Wuhan)
Wuhan, China
zhoufeng@cug.edu.cn

Javier Martin-Torres

University of Aberdeen, School of Geosciences
Aberdeen, United Kingdom
javier.martin-torres@abdn.ac.uk

Nectaria Diamanti

Aristotle University of Thessaloniki, Faculty of Sciences, Department of Geophysics
Thessaloniki, Greece
ndiamant@geo.auth.gr

ABSTRACT

Ground-penetrating radar (GPR) is a mature geophysical method that has gained increasing popularity in planetary science over the past decade. GPR has been utilised both for Lunar and Martian missions providing pivotal information regarding the near surface geology of Terrestrial planets. Within that context, numerous processing pipelines have been suggested to address the unique challenges present in planetary setups. These processing pipelines often require manual tuning resulting to ambiguous outputs open to non-unique interpretations. These pitfalls combined with the large number of planetary GPR data (kilometers in magnitude), highlight the necessity for automatic, objective and advanced processing and interpretation schemes. The current paper investigates the potential of deep learning for interpreting and processing GPR data. The one-shot multi-offset configuration is investigated via a coherent numerical case study, showcasing the potential of deep learning for A) reconstructing the dielectric distribution of the the near surface of Terrestrial planets, and B) filling missing or bad-quality traces. Special care was taken for the numerical data to be both realistic and challenging. Moreover, the generated synthetic data are properly labelled and made publicly available for training future data-driven pipelines and contributing towards developing pre-trained foundation models for GPR.

Keywords Deep Learning · Machine Learning · Chang'E-4 · Yutu-2 · Ground Penetrating Radar · GPR · Perseverance · inversion · FWI · big data · data imputation

1 Introduction

In-situ ground-penetrating radar (GPR) was introduced in planetary science in 2013 as part of the scientific payloads of Yutu-1 the Lunar rover of the Chinese Chang'E-3 mission [1]. Since then, GPR was part of the scientific payloads of the Yutu-2 rover from the Chang'E-4 mission [2], the lander of the Chang'E-5 [3] and Chang'E-6 missions, the rover Perseverance [4] from Mars 2020, and Zhurong rover from Tianwen-1 mission [5]. Moreover, GPR is planned to be used in the future missions Chang'E-7 [6] and ExoMars [7], both missions expected to take place before 2030. Consequently, GPR is one of the most mainstream scientific payloads in the new era of space exploration, and one of the few in-situ geophysical methods applied in planetary science.

GPR is a mature geophysical method with a wide range of applications, each with its own uniqueness and challenges [8]. To that extent, bespoke processing, modelling and interpretation tools are needed for almost each unique GPR application. Regarding planetary GPR, numerous processing and interpretation methodologies have been suggested, from hyperbola fitting [9, 10, 11] and migration [11] to centroid frequency shift analysis [11, 12, 13]. These approaches provide non-unique interpretations and often require laborious and manual processing combined with high-level of domain knowledge in GPR and geophysics in general.

The issues mentioned in the previous paragraph are not exclusive to planetary GPR. To address them, machine learning (ML) and data-driven approaches have been proposed in an attempt to develop objective and fully automatic interpretation frameworks primarily for Terrestrial applications. Fine-tuned foundation computer vision models and self-supervised learning have been used for automatic hyperbola fitting for non-destructive testing [14, 15]; convolutional neural networks and dense neural networks were suggested for removing background and ringing noise [16, 17, 18, 19]; random forest was used to estimate the diameters of rebars in concrete inspections [20]; convolutional auto-encoders were utilised for full waveform inversion (FWI) trained using synthetic data [21]; U-net was used to fill missing traces [22]; and unsupervised clustering was used to segment radagrams from Yutu-2 rover [23].

In the current paper we want to investigate the potential of machine learning for processing and interpreting one-shot multi-offset GPR data collected in a planetary setup [24]. In particular, we will conduct a coherent numerical study using realistic and challenging synthetic data focusing on two problems A) automatic full waveform inversion and B) filling missing or bad quality traces. The models will be tuned for frequencies around 60 – 100 MHz, similar to the range of the low frequency systems employed in the planetary rovers Yutu-1, Yutu-2 and Zhurong. Moreover our investigation will focus on media with low electric permittivity and conductivity as expected in the Lunar [25] and Martian subsurface [26].

2 One-Shot Multi-Offset Configuration

The most common measurement configuration in planetary in-situ GPR is the common-offset (CO), where the transmitter and receiver are moving along the measurement line while keeping their distance constant. CO has numerous drawbacks, one of the most important is the lack of a reliable methodology for calculating the subsurface permittivity. Hyperbola fitting (HF) is a very common approach used in CO-GPR for estimating the permittivity, however HF is reliable only for shallow layers [11], since for deeper ones' large boulders are necessary to give rise to visible hyperbolic reflectors; and the interpretation has a high rate of uncertainty with increasing depth [27].

The lack of an attainable method to calculate the permittivity in CO-GPR can be tackled using multi-offset (MO) configurations. where the receiver stands still while the transmitter is moving (or vice versa). MO-GPR has the ability to detect layers and accurately estimate their permittivity using the normal move-out (NMO) and Dix conversion [24, 28]. In a planetary setup, a MO-GPR can be realised –in theory– by placing a receiver on the lander of the mission [24]. The Chinese missions Chang'E-3, E-4 and Tianwen-1, all have landers, that could –in principle– be utilised for MO-GPR. Figure 1 illustrates a numerical case study using one-shot MO-GPR for mapping Lunar ejecta. The processing applied was a typical NMO and Dix conversion [24, 28]. The discretisation of the model is 2 cm, and the central frequency of the source is 80 MHz. It is apparent that via this approach deep layers can be sufficiently detected and characterised.

3 Deep Learning - Processing and Interpretation

In this section we will explore the capabilities of deep learning for processing and interpreting one-shot MO-GPR data. We will explore two problems A) Full-Waveform Inversion (FWI), and B) imputing missing traces. The section is divided into three subsections. The first subsection describes the synthetic training data used for training, providing details on how to generate them using the open-source electromagnetic simulator gprMax [29]. The second subsection suggests a deep learning scheme for automatic FWI using one-shot MO-GPR data. Finally, the last subsection investigates the capabilities of deep learning on reconstructing missing or bad quality traces.

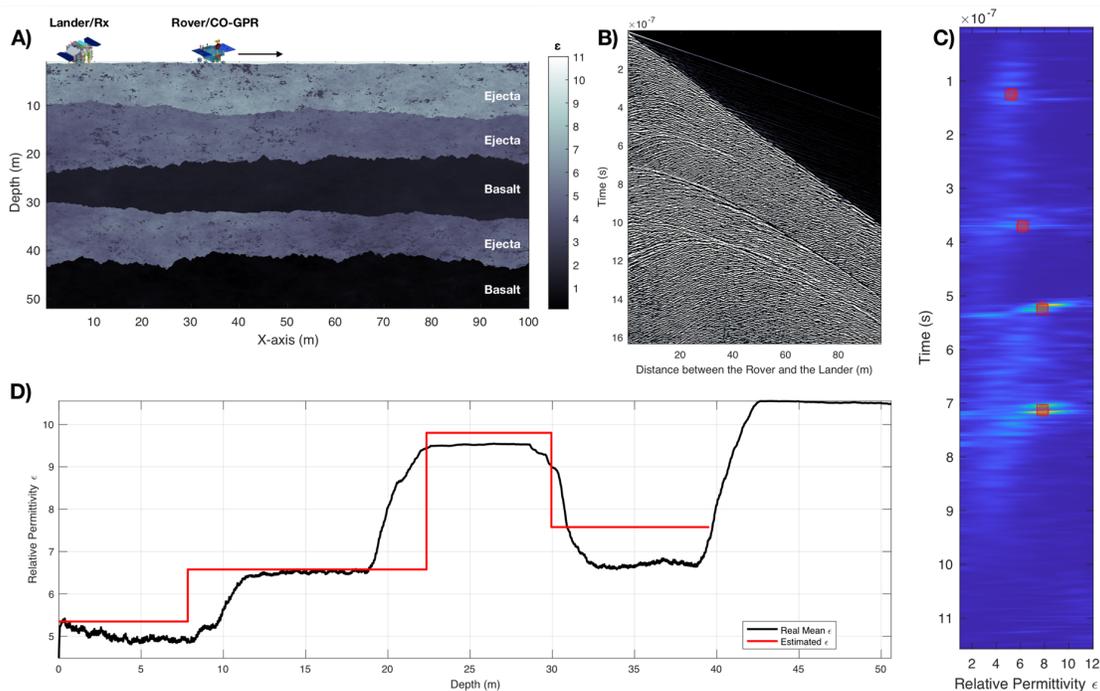


Figure 1: A) A generic one-shot multi-offset numerical case study. The discretisation of the model is 2 cm, and the central frequency of the source is 80 MHz. B) Processed B-Scan. C) NMO for picking velocities. D) The mean permittivity profile (black) and the reconstructed one (red) using Dix conversion [24].

3.1 Training Data

The synthetic training data were generated using the open-source software *gprMax* [29], an electromagnetic simulator that uses a second order accuracy (in both space and time) Finite-Differences Time-Domain (FDTD) method [30]. *GprMax* is capable of running in parallel using multiple CPUs and GPUs [31] making it an appealing choice for generating big data. Moreover, *gprMax* is tuned for GPR modelling, and is also script-able allowing for the generation of randomly varying models.

All the models of the training data are 2D with a uniform discretization step $\Delta x = \Delta y = 2$ cm. The time-step equals with $dt = 0.0467$ ns, which is 0.99 times the Courant limit. FDTD runs for 15000 iterations i.e. for a total of 700 ns. The dimensions of the model are 50.02×26 m, which results to a 2501×1300 2D grid. Perfectly Matched Layer (PML) is used to truncate the boundaries. In particular, a 60-layer thick PML is used in order to reduce the numerical artifacts from the surface waves on the free-space/soil interface [32]. The excitation source is placed 2 meters from the left boundary, and at 60-100 cm height from the surface. Both the shape and the central frequency of the pulse varies between training samples. The central frequency varies from 60-100 MHz, and the type of the pulse is randomly selected from *gaussiantdot*, *gaussiantdotnorm*, *gaussiantdotdot* and *ricker* (see *gprMax* documentation). Using different central frequencies and different types of pulses will allow for the model to generalise for unknown excitation sources. The polarisation of the pulse is orthogonal to the 2D domain. Along the surface, on the same height with the transmitter are 230 receivers, placed every 20 cm starting at a 2 meters distance from the excitation source.

Each of the models consists of a series of layers with random thickness, position, topography and dielectric properties. Moreover, within each layer the permittivity varies stochastically using fractal correlated noise with randomly selected fractal dimension [33]. The permittivity of the layers varies from $\epsilon = 2 - 10$, which is consistent with the permittivity values measured in the Lunar samples from the Apollo mission [25]. The conductivity in planetary setups is negligible (due to the absence of liquid water), and is primarily related to oxides of titanium and iron (e.g. ilmenite) [34]. For simplicity, we assume a correlation between permittivity and conductivity $\sigma = \epsilon/1000$, following the rationale that the

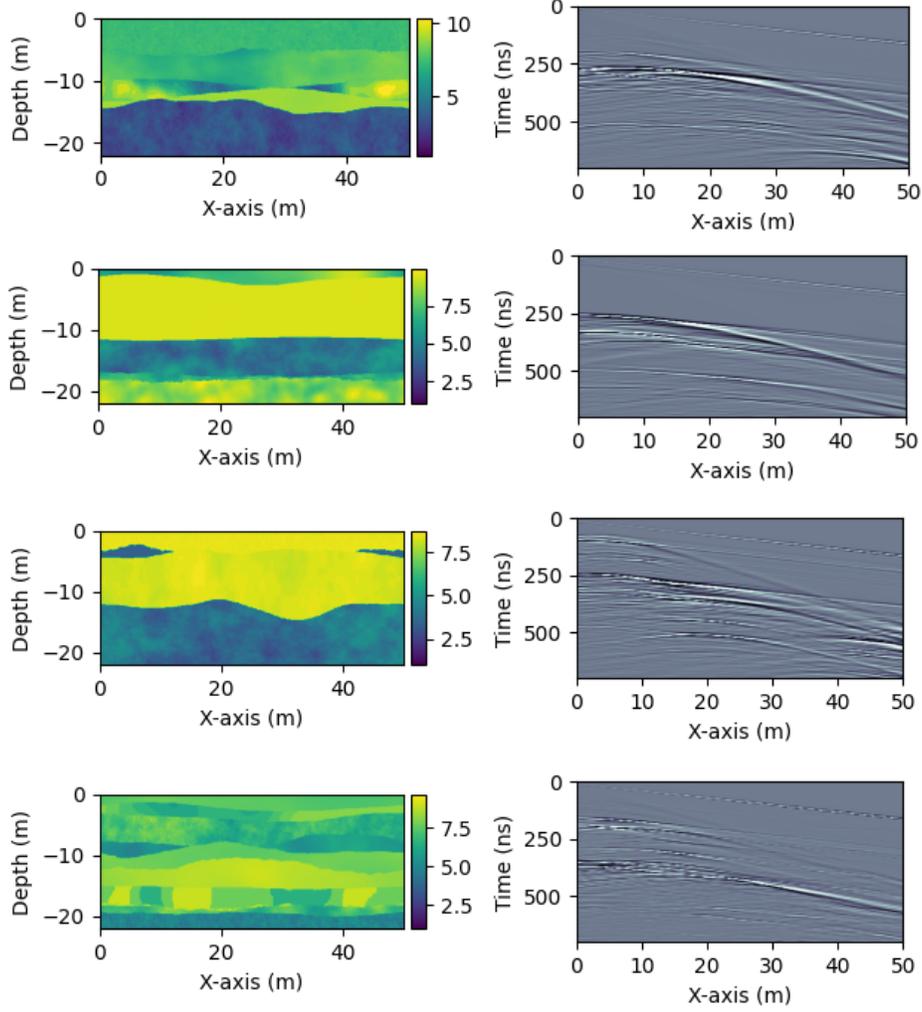


Figure 2: A set of training data from the 4400 samples used in this study. The colorbar illustrates the relative electric permittivity. Left figures depict the numerical models, and right figures are their corresponding B-Scans. A time-varying gain and signal saturation is applied to the B-Scans for illustration purposes.

abnormally high permittivity value of ilmenite and its dispersive properties will increase both the overall permittivity and the electromagnetic losses of an ilmenite-bearing formation [35].

The training data consists of 4400 samples. Figure 2 shows 4 samples and their corresponding B-Scans. Varying layers, with different geometries and dielectric properties are included in the training data to increase the generalization capabilities of the resulting machine learning models. The initial raw B-Scan $\mathbf{B}_i \in \mathbb{R}^{230 \times 15000}$ is a matrix with 230 traces each of which having 15000 points length, where i is the sample number. A time-varying gain is applied to every trace $A(n \cdot dt) = A(n \cdot dt) \cdot n^3$ to enhance late reflections, where n is the time index. Subsequently each trace is compressed and resized to 230 dimensions, resulting to a new time step that is equal to 3.0434 ns. The processed and compressed B-Scan is now expressed as $\mathbf{Q}_i \in \mathbb{R}^{230 \times 230}$, a square matrix with 230 by 230 dimensions.

The initial dielectric models can be expressed as matrices $\mathbf{M}_i \in \mathbb{R}^{2501 \times 1300}$ containing the permittivity values of the 50.02×26 m domain, with a spatial step of $\Delta x = \Delta y = 2$ cm. The first 4 meters of the model are free-space, and therefore are trimmed out i.e. $\mathbf{M}_i \in \mathbb{R}^{2501 \times 1100}$. Subsequently, a nearest neighbor interpolation is applied to reduce the dimensions of the data to $\mathbf{W}_i \in \mathbb{R}^{224 \times 224}$ resulting to the new spatial steps $\Delta x = 22.32$ cm and $\Delta y = 9.82$ cm.

Conclusively, the final training data consist of 4400 compressed and processed (time-varying gain) B-Scans $\mathbf{Q}_i \in \mathbb{R}^{230 \times 230}$, and their corresponding dielectric models $\mathbf{W}_i \in \mathbb{R}^{224 \times 224}$.

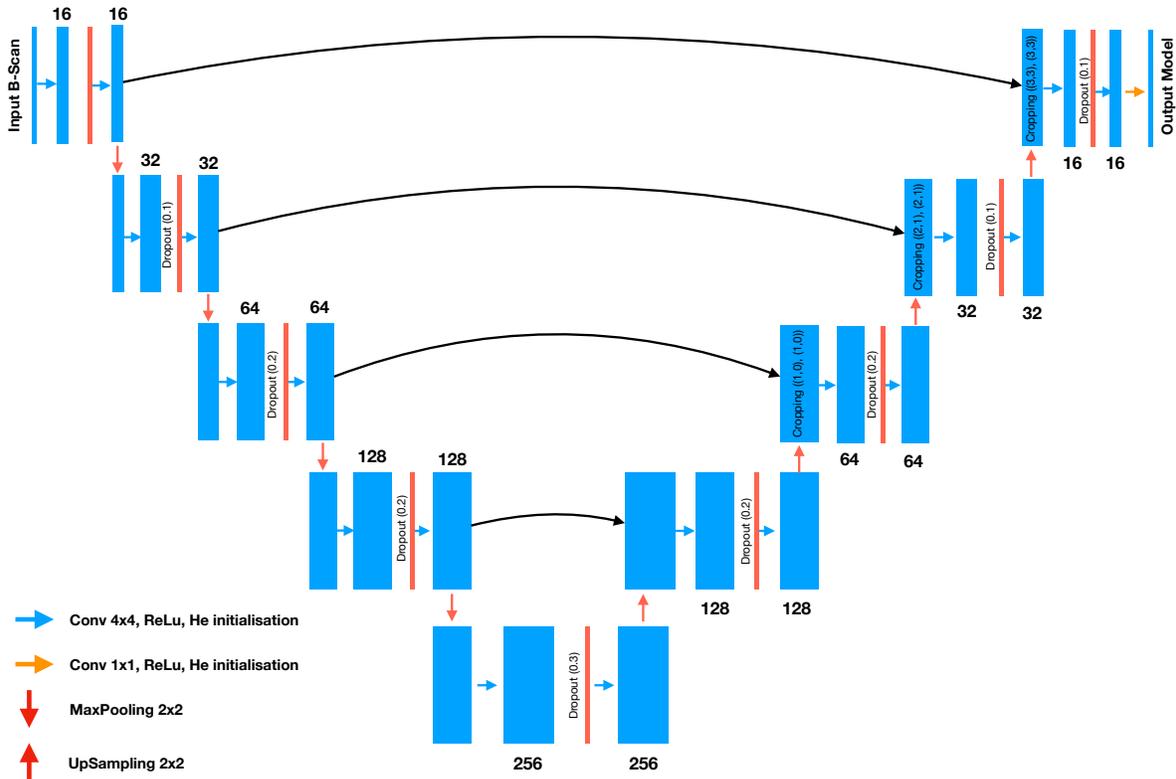


Figure 3: The architecture of the ensemble U-nets used for FWI. The input B-Scans have 230×230 dimensions each, and the output models have 224×224 .

3.2 Full-Waveform Inversion

In this section we aim to reconstruct the dielectric models \mathbf{W}_i from their corresponding B-Scans \mathbf{Q}_i . In an abstract manner this is an image-to-image translation where the input image (B-Scan) $\mathbf{W}_i \in \mathbb{R}^{230 \times 230}$ is mapped into another image (dielectric model) $\mathbf{Q}_i \in \mathbb{R}^{224 \times 224}$. A very common deep learning methodology for this type of problems is U-net [36]. More advanced methods suitable for image-to-image translation are Generative Adversarial Networks (GAN) [37], Pix2Pix [38]), Visual Transformers [39] etc.. Despite the benefits of these techniques, in this paper we will employ a simple U-net in order to use it as a benchmark for comparing more advanced methodologies in the future.

The architecture of the U-net is shown in Figure 3. Dropout layers are used both in the encoding and decoding parts. The weights were initialised using the Kaiming He method [40]. The Adam optimizer [41] was used to train the model. The learning rate was initially set to $5e - 4$, after 55 epochs the learning rate reduced to $1e - 4$, and after 30 epochs it was further reduced to $5e - 5$ for another 30 epochs. The batch size equals with 5, and the loss function is the mean absolute difference. A validation set (10%) is used to monitor the loss and stop training if the validation loss does not improve for 30 consecutive epochs.

The U-net architecture is shown in Figure 3. The training process described in the previous paragraph is executed 10 times using different initial weights, resulting to 10 different U-nets tuned for the synthetic training samples. The final model is an ensemble of the 10 U-nets. The performance of the model is evaluated in unknown scenarios that were not included in the training data. Figure 4 shows a set of examples of the FWI executed using the ensemble U-nets. It is apparent, that deep learning has the capability to reconstruct the dielectric properties of the subsurface even with just using one-shot MO-GPR data. More advanced approaches, combined with novel pre-processing techniques can further increase the overall accuracy, providing robust and reliable foundation models for FWI.

3.3 Fill Missing Data

Missing or bad quality data is a typical problem encountered both in Terrestrial and planetary applications of GPR [42, 43]. In this section we explore the capabilities of deep learning on reconstructing missing traces in one-shot

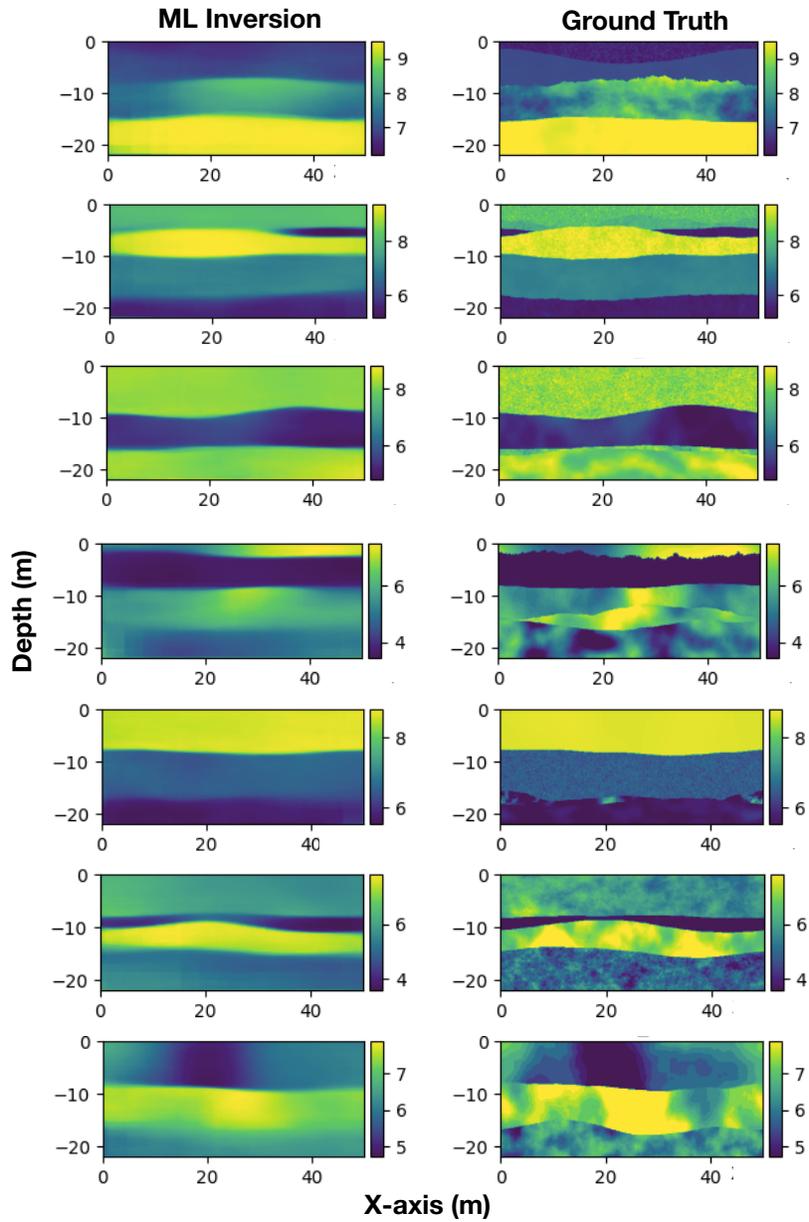


Figure 4: A set of examples comparing the ground truth permittivity models to the reconstructed ones using the suggested ensemble U-nets. These are unknown models that were not included in the training process. Colourbars depict the relative permittivity values. It is apparent that the proposed ensemble approach has the capability to reconstruct a smooth representation of the subsurface.

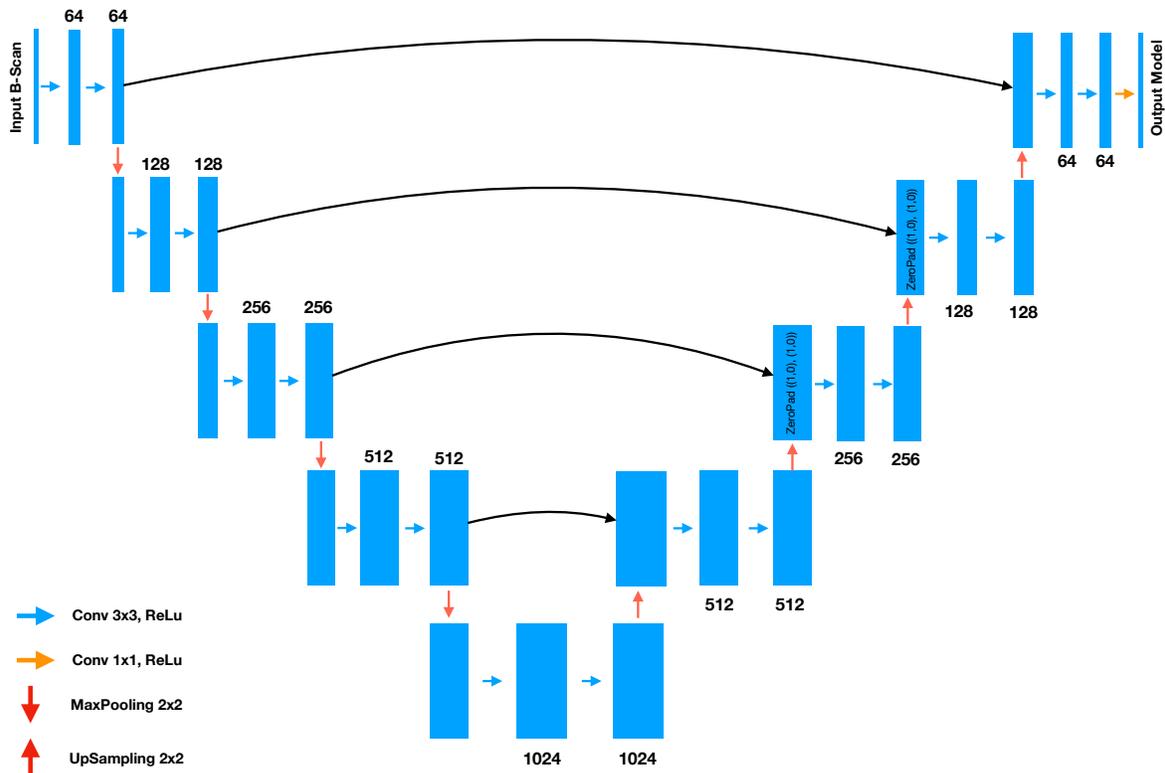


Figure 5: The architecture of the U-net used for filling missing data. The inputs are the corrupted B-Scans each having 230×230 dimensions, and the outputs are the initial B-Scans.

MO-GPR data. For the training we initially corrupt the data by removing traces from the original radagram. We keep one trace every 2 meters, i.e. for every 10 traces 9 traces are removed. The empty data are replaced with zeros. The corrupted data $\mathbf{D}_i \in \mathbb{R}^{230 \times 230}$ are now used as inputs, and the initial data $\mathbf{W}_i \in \mathbb{R}^{230 \times 230}$ are the desired outputs.

Similarly to the previous sub-section, U-net is used to map the causal relationship between \mathbf{D} and \mathbf{W} . U-net is a very typical scheme used to fill missing data in geophysics, with promising results in MO seismic data [44]. The architecture of the employed U-net is shown in Figure 5. The optimizer used is Adam with learning rate $1e - 4$. The batch size is 16, and the loss function is the mean squared error. Similarly to the previous section 10% of the data are used as a validation set with a patience of 30 epochs.

Figure 6 shows a series of examples demonstrating the capabilities of deep learning and U-net on reconstructing missing or bad quality data. Even in the absence of 90% of the data, the trained U-net was capable of sufficiently reconstructing and predicting the missing values.

4 Data Availability - Kaggle Competition

The training data used in this study are available in Kaggle in the community competition **GprMax Deep Learning Challenge 1 (GDLC-1)**. The competition aims at developing foundation deep learning models for FWI of one-shot MO-GPR data. Any submitted model must be trained subject to synthetic data (ideally using gprMax). The participants can either use the provided synthetic datasets, or they can use gprMax (or any software of their choice) to generate additional data to complement training. The participants can also apply any pre-processing they think necessary to the input files prior to training. The final evaluation metric is the mean absolute error between the ground truth permittivity of the testing models and the reconstructed ones using the proposed ML models.

Files - Training Data

- **Training_Bscan**: Folder that contains the .npy files for the processed (time-varying gain) B-Scans.
- **Bscan_'n'.npy**: Input data. The nth Bscan saved in a numpy array, $\text{np.shape}(\text{B_Scan_}'n'.\text{npy}) = (230, 230)$.

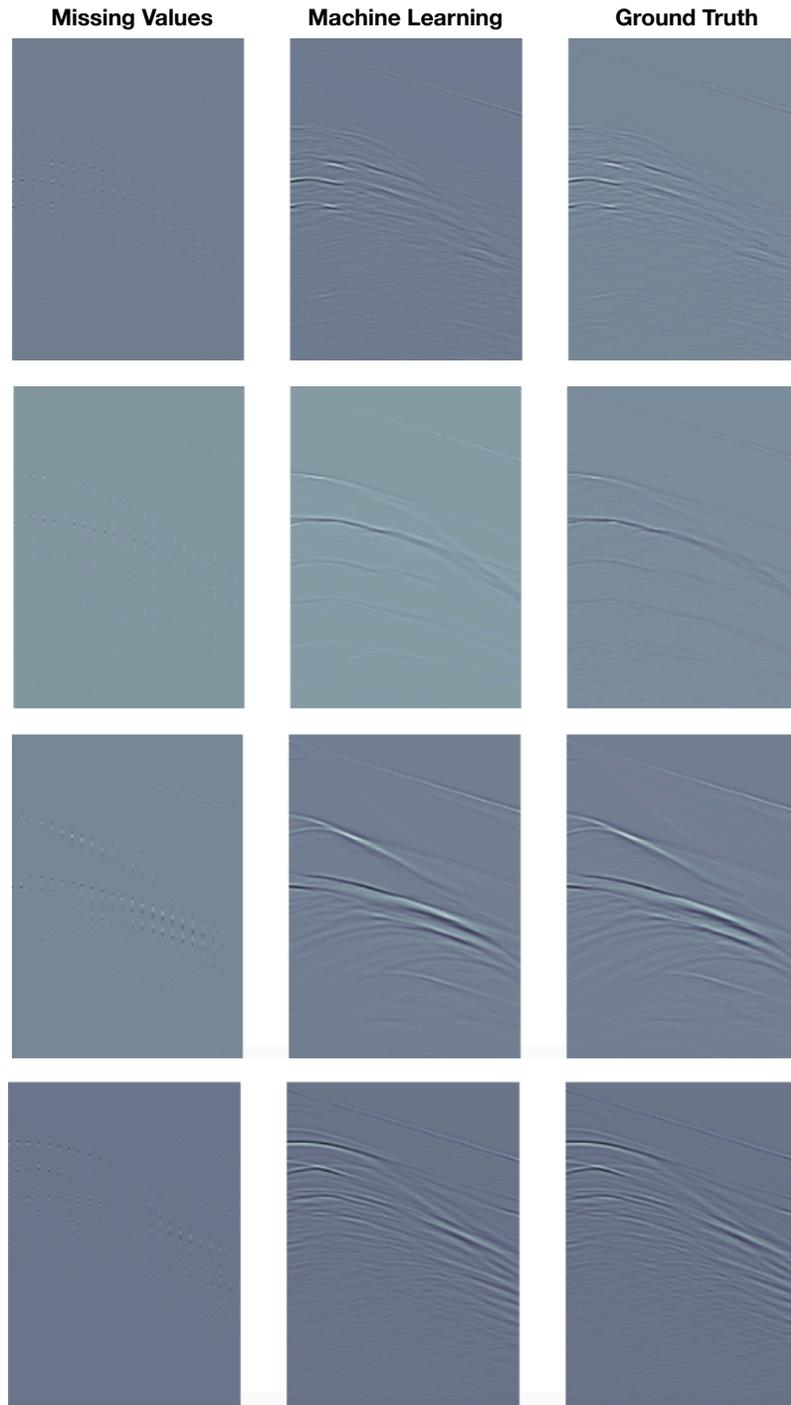


Figure 6: A set of testing examples used to evaluate the performance of U-net for reconstructing missing or bad quality MO-GPR data. These samples were not included in the training process. The label "Missing Values" correspond to the corrupted data with 9 traces removed every 2 meters. The label "Machine Learning" refers to the reconstructed B-Scans using the proposed U-net, and "Ground Truth" refers to the original B-Scans. It is apparent that deep learning is an appealing choice for treating missing values in one-shot MO GPR data.

- **Training_Labels:** Folder that contains the .npy files for the ground truth i.e. relative permittivity structure.
- **Model_`n`.npy:** Labels. The nth ground truth model saved in a numpy array, $\text{np.shape}(\text{Model_`n`.npy}) = (224,224)$.

Files - Testing Data

- **Evaluation_Dataset:** Folder that contains the .npy files for the processed (time-varying gain) B-Scans from the testing set.
- **Testing_Bscan_`n`.npy:** Input data. The nth testing Bscan saved in a numpy array, $\text{np.shape}(\text{B_Scan`n`.npy}) = (230,230)$.

The B-Scans in the training set are complete without missing traces, while the B-Scans in the evaluation set have some random missing set of traces. This is to add an extra challenge to the competition simulating missing data or bad quality data that need to be interpolated. The participants need to find a way to train their model for missing data or develop a way to fill the missing data prior to FWI.

5 Conclusions

A numerical case study was presented to demonstrate the potential capabilities of one-shot multi-offset measurement configuration combined with deep learning interpretation. Through a series of numerical tests we showcase that deep learning can sufficiently invert for the subsurface's dielectric properties, and accurately reconstruct missing or bad quality data. Despite trained with numerical data, these models can be used as foundation models for future transfer learning with real-data, and act as case studies for demonstrating the potential capabilities of different measurement configurations and machine learning models for planetary GPR.

References

- [1] Guang-You Fang, Bin Zhou, Yi-Cai Ji, Qun-Ying Zhang, Shao-Xiang Shen, Yu-Xi Li, Hong-Fei Guan, Chuan-Jun Tang, Yun-Ze Gao, Wei Lu, et al. Lunar penetrating radar onboard the chang'e-3 mission. *Research in astronomy and astrophysics*, 14(12):1607, 2014.
- [2] Chunlai Li, Yan Su, Elena Pettinelli, Shuguo Xing, Chunyu Ding, Jianjun Liu, Xin Ren, Sebastian E. Lauro, Francesco Soldovieri, Xingguo Zeng, Xingye Gao, Wangli Chen, Shun Dai, Dawei Liu, Guangliang Zhang, Wei Zuo, Weibin Wen, Zhoubin Zhang, Xiaoxia Zhang, and Hongbo Zhang. The moon's farside shallow subsurface structure unveiled by Chang'E-4 lunar penetrating radar. *Science Advances*, 6(9), 2020.
- [3] Feiyang Fang, Chunyu Ding, Jianqing Feng, Yan Su, Ravi Sharma, and Iraklis Giannakis. In-situ radar observation of shallow lunar regolith at the chang'e-5 landing site: Research progress and perspectives. *Remote Sensing*, 15(21), 2023.
- [4] Svein-Erik Hamran, David A. Paige, Hans E. F. Amundsen, Tor Berger, Sverre Brovoll, Lynn Carter, Leif Damsgård, Henning Dypvik, Jo Eide, Sigurd Eide, Rebecca Ghent, Øystein Hellereen, Jack Kohler, Mike Mellon, Daniel C. Nunes, Dirk Plettmeier, Kathryn Rowe, Patrick Russell, and Mats Jørgen Øyan. Radar imager for mars' subsurface experiment—rimfax. *Space Science Reviews*, 216:128, 2020.
- [5] Lei Zhang, Chao Li, Jinhai Zhang, Bin Zhou, Yu-Yan Sara Zhao, Yang Liu, Kaichang Di, Ross N. Mitchell, Juan Li, Zhigang Zhang, Lin Chen, Xiaofeng Liang, Weijia Sun, Yike Liu, Xu Zhao, Jinlai Hao, Changyi Xu, Yikang Zheng, Yibo Wang, Xin Wang, Pan Zhao, Wenmin Lv, Yang Li, Haiqiang Lan, and Yongxin Pan. Buried palaeo-polygonal terrain detected underneath utopia planitia on mars by the zhurong radar. *Nature Astronomy*, 2023.
- [6] Yongliao Zou, Yang Liu, and Yingzhuo Jia. Overview of china's upcoming chang'e series and teh scientific objectives and payloads for chang'e-7 mission. *51st Lunar and Planetary Science Conference*, 2020.
- [7] Y. Hervé, V. Ciarletti, A. Le Gall, C. Corbel, R. Hassen-Khodja, W.S. Benedix, D. Plettmeier, O. Humeau, A.J. Vieau, B. Lustremont, S. Abbaki, E. Bertran, L. Lapauw, V. Tranier, N. Oudart, F. Vivat, C. Statz, Y. Lu, S. Hegler, and A. Hérique. The wisdom radar on board the exomars 2022 rover: Characterization and calibration of the flight model. *Planetary and Space Science*, 189:104939, 2020.
- [8] David J. Daniels. *Ground Penetrating Radar*. U.K.:Institution of Engineering and Technology, 2004.
- [9] Z. Dong, G. Fang, D. Zhao, B. Zhou, Y. Gao, and Y. Ji. Dielectric properties of lunar subsurface materials. *Geophysical Research Letters*, 47, 2020.

- [10] Feng, Jianqing, Siegler, Matthew. A., and White, Mackenzie N. Dielectric properties and stratigraphy of regolith in the lunar south pole-aitken basin: Observations from the lunar penetrating radar. *A&A*, 661:A47, 2022.
- [11] Iraklis Giannakis, Javier Martin-Torres, Maria-Paz Zorzano, Yan Su, Craig Warren, and Antonios Giannopoulos. Stochastic hyperbola fitting, probabilistic inversion, reverse-time migration and clustering: A novel interpretation toolbox for in-situ planetary radar. *Icarus*, 400:115555, 2023.
- [12] Chunyu Ding, Zhiyong Xiao, Yan Su, Cheng Tang, Bin Liu, Yunlong Wu, and Xuejun Zhang. Compositional variations along the route of chang’e-3 yutu rover revealed by the lunar penetrating radar. *Progress in Earth and Planetary Science*, 7(1):1–16, 2020.
- [13] Sigurd Eide, Titus M. Casademont, Tor Berger, Henning Dypvik, Emileigh S. Shoemaker, and Svein-Erik Hamran. Radar attenuation in the shallow martian subsurface: Rimfax time-frequency analysis and constant-q characterization over jezero crater floor. *Geophysical Research Letters*, 50(7):e2022GL101429, 2023.
- [14] Jian Huang, Xi Yang, Feng Zhou, Xiaofeng Li, Bin Zhou, Song Lu, Sergey Ivashov, Iraklis Giannakis, Fannian Kong, and Evert Slob. A deep learning framework based on improved self-supervised learning for ground-penetrating radar tunnel lining inspection. *Computer-Aided Civil and Infrastructure Engineering*, 39(6):814–833, 2024.
- [15] Xiaofeng Li, Hai Liu, Feng Zhou, Zhongchang Chen, Iraklis Giannakis, and Evert Slob. Deep learning-based nondestructive evaluation of reinforcement bars using ground-penetrating radar and electromagnetic induction data. *Computer-Aided Civil and Infrastructure Engineering*, 37(14):1834–1853, 2022.
- [16] Hai-Han Sun, Weixia Cheng, and Zheng Fan. Clutter removal in ground-penetrating radar images using deep neural networks. In *2022 International Symposium on Antennas and Propagation (ISAP)*, pages 17–18, 2022.
- [17] Hai-Han Sun, Weixia Cheng, and Zheng Fan. Learning to remove clutter in real-world gpr images using hybrid data. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–14, 2022.
- [18] Ourania Patsia, Antonios Giannopoulos, and Iraklis Giannakis. Background removal, velocity estimation, and reverse-time migration: A complete gpr processing pipeline based on machine learning. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–11, 2023.
- [19] Ourania Patsia, Antonios Giannopoulos, and Iraklis Giannakis. Gpr full-waveform inversion with deep-learning forward modeling: A case study from non-destructive testing. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–10, 2023.
- [20] Iraklis Giannakis, Antonios Giannopoulos, and Craig Warren. A machine learning scheme for estimating the diameter of reinforcing bars using ground penetrating radar. *IEEE Geoscience and Remote Sensing Letters*, 18(3):461–465, 2021.
- [21] Zi Xian Leong and Tiejuan Zhu. Direct velocity inversion of ground penetrating radar data using gprnet. *Journal of Geophysical Research: Solid Earth*, 126(6):e2020JB021047, 2021.
- [22] Qiqi Dai, Yee Hui Lee, Hai-Han Sun, Genevieve Ow, Mohamed Lokman Mohd Yusof, and Abdulkadir C. Yucel. 3dinvnet: A deep learning-based 3d ground-penetrating radar data inversion. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–16, 2023.
- [23] Iraklis Giannakis, Ciaran McDonald, Jianqing Feng, Feng Zhou, Yan Su, Javier Martin-Torres, Maria-Paz Zorzano, Craig Warren, Antonios Giannopoulos, and Georgios Leontidis. Automatic segmentation of radar data from the chang’e-4 mission using unsupervised machine learning: A data-driven interpretation approach. *Icarus*, 417:116108, 2024.
- [24] I. Giannakis, J. Martin-Torres, M. P. Zorzano, C. Warren, and A. Giannopoulos. Low frequency multi-offset planetary radar: A novel mission scheme that utilizes both rover and lander. In *AGU Fall Meeting Abstracts*, pages P55E–1611, 2022.
- [25] G. R. Olhoeft and D. W. Strangway. Dielectric properties of the first 100 meters of the moon. *Lunar Sourcebook*, 24:394–404, 1975.
- [26] Ling Zhang, Yi Xu, Renrui Liu, Ruonan Chen, Roberto Bugiolacchi, and Rui Gao. The dielectric properties of martian regolith at the tianwen-1 landing site. *Geophysical Research Letters*, 50(13):e2022GL102207, 2023.
- [27] Iraklis Giannakis, Feng Zhou, Craig Warren, and Antonios Giannopoulos. On the limitations of hyperbola fitting for estimating the radius of cylindrical targets in nondestructive testing and utility detection. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2022.
- [28] Dimitrios Angelis, Craig Warren, Nectaria Diamanti, James Martin, and A. Peter Annan. Development of a workflow for processing ground-penetrating radar data from multiconcurrent receivers. *GEOPHYSICS*, 87:WB9–WB18, 2022.

- [29] C. Warren, A. Giannopoulos, and I. Giannakis. gprMax: Open source software to simulate electromagnetic wave propagation for ground penetrating radar. *Computer Physics Communications*, 209:163–170, 2016.
- [30] A. Taflove and S. C. Hagness. *Computational electrodynamics: The finite-difference time-domain method*. Norwood, MA, USA: Artech House, 2000.
- [31] Craig Warren, Antonios Giannopoulos, Alan Gray, Iraklis Giannakis, Alan Patterson, Laura Wetter, and Andre Hamrah. A CUDA-based GPU engine for gprMax: Open source FDTD electromagnetic simulation software. *Computer Physics Communications*, 237:208 – 218, 2019.
- [32] A. Giannopoulos. An improved new implementation of complex frequency shifted PML for the FDTD method. *IEEE Transactions on Antennas and Propagation*, 56(9):2995–3000, 2008.
- [33] I. Giannakis, A. Giannopoulos, and C. Warren. A realistic FDTD numerical modeling framework of ground penetrating radar for landmine detection. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 9(1):37–51, 2016.
- [34] D. W. Strangway, G. W. Pearce, and G. R. Olhoeft. Magnetic and dielectric properties of lunar samples. In *The Soviet-Am. Conf. on Cosmochem. of the Moon and Planets*, pages 417–431, Washington, 1977. NASA.
- [35] Iraklis Giannakis, Javier Martin-Torres, Yan Su, Jianqing Feng, Feng Zhou, Maria-Paz Zorzano, Craig Warren, and Antonios Giannopoulos. Evidence of shallow basaltic lava layers in von kármán crater from yutu-2 lunar penetrating radar. *Icarus*, 408:115837, 2024.
- [36] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015.
- [37] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014.
- [38] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks, 2018.
- [39] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021.
- [40] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, 2015.
- [41] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- [42] Deniz Kumlu, Kubra Tas, and Isin Erer. Missing data recovery via deep networks for limited ground penetrating radar measurements. *Remote Sensing*, 14(3), 2022.
- [43] Chongmin Zhang, Mathieu Gravey, Grégoire Mariéthoz, and James Irving. Reconstruction of missing gpr data using multiple-point statistical simulation. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–17, 2024.
- [44] Xiaoning Chai, Haixia Gu, Fang Li, et al. Deep learning for irregularly and regularly missing data reconstruction. *Scientific Reports*, 10:3302, 2020.