

MODELING GPR OBSERVATIONS ON RAILWAY TRACKS VIA BLACK BOX AND PHYSICS INFORMED NEURAL NETWORKS

THOMAS RIGONI¹, GIACOMO ARCIERI², MARCUS
HAYWOOD-ALEXANDER², DAVID HAENER³, ELENI CHATZI²

¹ Department of Computer Science
ETH Zürich
Universitätstrasse 6, 8092 Zürich, Switzerland
e-mail: thomas.rigoni@usys.ethz.ch thomas.rigoni7@gmail.com

² Institute of Structural Engineering (IBK)
ETH Zürich
Stefano-Franscini-Platz 5, 8093 Zürich, Switzerland
e-mail: {giacomo.arciери, marcus.haywood-alexander, chatzi}@ibk.baug.ethz.ch

³ Swiss Federal Railways (SBB)
e-mail: david.haener@sbb.ch

Key words: Ground Penetrating Radar (GPR), Physics-Informed Neural Networks (PINNs), Structural Health Monitoring (SHM), Railway Infrastructure

Summary. The assessment of track condition is important for the upkeep of modern rail infrastructure networks. Current methods for measuring and tracking the deterioration of the ballast layer are invasive, costly, and time-consuming, posing challenges which reduce the ability to adopt preventive maintenance strategies. Recent research investigates the utilization of Ground Penetrating Radar (GPR) as a cost-effective, fast, and non-invasive alternative for estimating ballast fouling and sub-surface water content. Radar recordings are data rich, resulting in spatially dense information in terms of the composition of the scanned substrate. While Machine Learning methods can in principle be used to automatically detect deterioration features, this would require high amounts of (labelled) training data, such that cannot be extracted in the field. Thus, simulation remains the main resort for ensuring such datasets. However, the generation of large-scale datasets using simulating software, such as gprMax, demands impractical amounts of computing power. In response to this challenge, we create two large simulated GPR datasets, which are respectively used to investigate the application of i) a CNN-based encoder-decoder architecture and ii) a physics-informed neural network (PINN) approach, to build an accurate surrogate model for gprMax. The former results in a successful surrogate, for which a reduction in computational cost of ≈ 300000 times is reported in a batched inference setting compared with the gprMax software, which results in a general speedup of the dataset generation process of ≈ 100 times. Our more efficient solution can be adopted for the creation of extensive datasets, paving the way for the subsequent monitoring methods aimed at supporting the maintenance of railway infrastructure. The PINN approach is tested as a promising scheme for exploiting Maxwell's equations to approximate the underlying electromagnetic wave field propagation. Ini-

tial experiments show that a straightforward implementation of PINNs for wave propagation in layered media faces significant challenges in this setting and more research in this direction is needed.

1 INTRODUCTION

Condition monitoring of railway infrastructure is imperative for efficient maintenance planning, ensuring high-quality service, extended lifespan, and enhanced safety of railway systems [1]. Current methods for assessing the condition of ballast and sub-ballast layers typically involve invasive, costly, and time-consuming procedures, such as excavation, which impede the adoption of proactive maintenance strategies [2]. Recent advancements propose the use of ground-penetrating radar (GPR) as a non-invasive, cost-effective, and rapid alternative for estimating ballast fouling and sub-surface moisture accumulation [3, 4], which comprise primary safety risks in railway infrastructure [5, 6]. Adoption of GPR technology could significantly accelerate track condition assessment, leading to comprehensive railway network inspection and timely maintenance interventions. Machine learning (ML) techniques can be applied for automatic detection of ballast fouling and groundwater presence in GPR imagery [7], but this requires extensive labeled training data, which is impractical to obtain through real-world excavation. To overcome this limitation, computer simulations can generate large GPR datasets with precise geometry maps and labels. However, creating large-scale datasets using suitable simulation tools, such as *gprMax*, is computationally demanding, particularly for realistic 3D geometries [8].

This study explores the development of deep learning-based surrogate models for *gprMax*, enabling the generation of large-scale labeled GPR datasets of railway track configurations. We investigate two approaches: a black-box convolutional neural network (CNN)-based encoder-decoder model and the application of physics-informed neural networks (PINNs). We opt for PINNs due to their ability to integrate constraints from partial differential equations (PDEs) in their training process [9, 10], specifically leveraging Maxwell’s equations to model the electromagnetic (EM) wave field propagation. PINNs have already been used to model wave propagation [11, 12], but no previous work has applied PINNs to GPR data with complexity comparable to railway track infrastructure. The main contributions of this work are: the creation of two simulated, realistic railway track GPR datasets, adaption and training of a black-box CNN model for surrogate modelling of GPR recordings, and exploration of PINN architectures to model EM wave field propagation, as a proposition for an alternative approach to GPR surrogate modelling.

2 PREVIOUS WORK

The software library *gprMax* [8] is widely used for GPR simulations due to its flexibility and computational efficiency, owing to its CUDA implementation. For example, Koyan et al. [13] created realistic 3D sedimentary models and corresponding GPR data using *gprMax*, while Giannakis et al. [14, 15] applied it for landmine detection. The use of ML for the modelling of GPR data has received limited attention. Giannakis et al. [16] used a multi-layer perceptron (MLP) to predict mean A-scan traces from a soil model described by five input parameters. In another study [17], they combined MLP with principal component analysis to estimate A-scan responses from a metallic rebar embedded in a concrete foundation. Dai et al. [18] employed an encoder-decoder architecture with attention-based feature fusion to compute the B-scan response of geometries using permittivity and permeability maps. Although effective, these frameworks

are still limited to simple geometries with single embedded objects surrounded by near-uniform materials.

Most research on PINNs for wave propagation utilizes the general scalar wave equation or its acoustic version. Moseley et al. [11] used a 10-layer PINN to solve the acoustic wave equation for various velocity models, demonstrating efficient time-domain simulation extensions, although with some inaccuracy in reflected wavefronts for complex models. Rasht-Behesht et al. [12] conducted experiments on acoustic wave propagation for seismic data, achieving precise forward simulations and Full Waveform Inversion (FWI). Their approach introduces free-surface boundary conditions to approximate planetary surfaces, which results in precise reflected waves. The application of PINNs to EM waves, and in particular to GPR data, has received less attention. Zheng and Wang [19] used PINNs to perform predictions beyond the scope of observations (time domain extension), showing potential to surpass Finite Difference Time Domain (FDTD) solvers in addressing numerical dispersion. They used an MLP with spatial and temporal coordinates as inputs to predict the electric field at specific points. However, the geometries studied were simple and the models struggled to capture high-frequency EM wave components. No experiment showcases a learned strong reflection of the EM waves in domain extension, with restriction of the predictive capacity within the observation domain. The authors noted that training times for their models were multiple hours, while FDTD solvers required only seconds for 2D solutions.

3 DATASET GENERATION

Two realistic and varied GPR railway track datasets are created, each with specific characteristics and use cases. Both datasets are intended for use with supervised ML algorithms. Each sample is composed of an input geometry map and the associated output GPR observation, while one dataset (A-scan outputs) also additionally includes EM field snapshots. The geometry maps are 2D matrices containing the EM properties of the sample, namely the relative permittivity ϵ_r , electrical conductivity σ , relative permeability μ_r , and magnetic loss σ_* . Output observations are either A-scan or B-scan simulation results generated via the open source library gprMax.

Both datasets implement three different railway track topologies: AC rail, regular gravel-sand subgrade, and a more compacted version of the latter, termed PSS in Switzerland. The GPR antenna is represented by a Hertzian dipole source emitting a Ricker wavelet at a central frequency of 1 GHz. Rails are not present in the dataset as the samples are in two dimensions. Three sleeper configurations are available, namely wooden, concrete, and steel sleepers. Ballast stones are represented by circles of varying diameter. Ballast compaction is performed by means of a 2D physics simulation to enable a more realistic placement of the stones. Ballast fouling is simulated in two modes: i) via the progressive reduction of the average diameter of ballast stones due to their progressive shredding and ii) via the addition of a background layer with varying height to simulate a higher fouling level. The subgrade and subsoil layers are represented by Peplinski mixed models of soil implemented in gprMax. All the layers present randomized depths and rough surfaces to better align to a real scenario. Diverse layer moisture contents are simulated by manipulating their respective water fraction. Groundwater infiltration is also reproduced, in which a water pocket is present between two layers. All assumed configurations, conditions, and deterioration levels follow probability distributions which reflect real-world data provided by the Swiss Federal Railways (SBB).

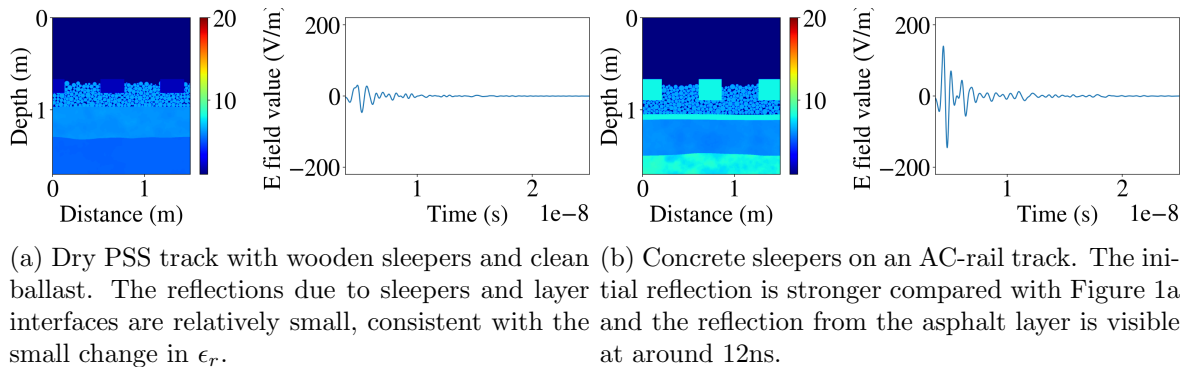


Figure 1: Relative permittivity plots (left) and central A-scans (right) of samples generated in the A-scan dataset.

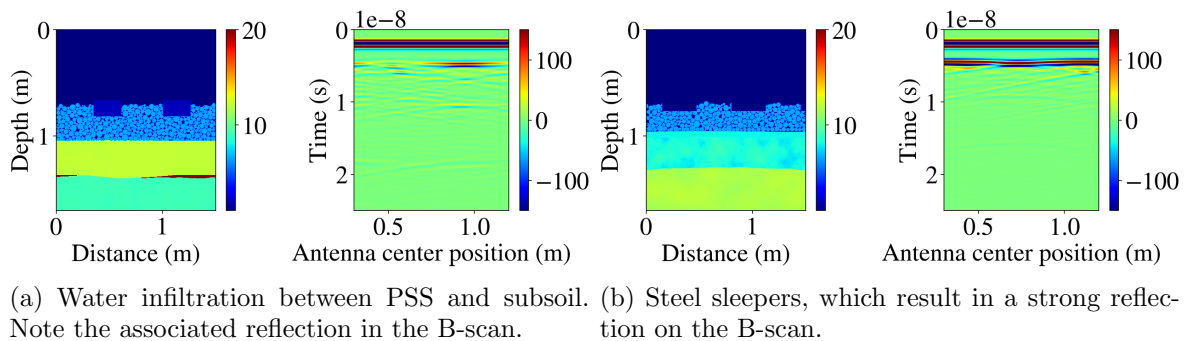


Figure 2: Relative permittivity plots (left) and resulting B-scans (right) of samples generated in the B-scan dataset.

The first dataset is composed by 30000 A-scan samples, where the antenna is placed centrally in the scene. Snapshots of electric and magnetic fields are produced at 1ns intervals. This dataset is suited to training of PINN models. The second dataset is formed by 4000 B-scan samples, each sample composed by 90 A-scans with a separation of 1cm between them. No snapshots of EM fields are recorded in this case. Figures 1 and 2 show relative permittivity maps and A/B-scans for samples in the two datasets.

4 BLACK BOX CNN MODEL

We present a black box CNN model adapted from Dai et al. [18]. This architecture consists of two separate convolutional encoders for relative permittivity ϵ_r and conductivity σ maps, a feature fusion module based on attention mechanisms, and a convolutional decoder which outputs the predicted B-scans. The model is here extended with i) an additional encoding block in both feature encoders, and ii) the use of a 3×3 convolution with 3×3 strides in the connection block responsible for the fusion of the attention-captured features, which effectively acts as a parametrized downsampling operation. This architecture was trained for 300 epochs with a batch size of 10. The Adam optimizer was used to minimise the mean squared error (MSE) loss between observations and output. The initial learning rate was set to 10^{-4} and a

LR scheduler was adopted. The best model was selected as the one with lowest validation loss. Some pre-processing of the input data was performed, namely i) resizing of the labels to the prediction size of 192×224 pixels, and ii) the pixel-wise median value of the labels was calculated and subtracted from the images. This latter step allows to remove the direct source-receiver wave from the radargrams and improves the empirical accuracy of the models.

Figure 3 shows some validation dataset geometry maps, ground truth B-scan, and predictions of the network. In particular, the model predictions show a remarkable similarity to the labels, with only negligible differences. To the best of our knowledge, no other published work shows predictions of GPR data with a comparable level of complexity of the input geometries and accuracy of the predicted B-scans, which form a salient contribution of this work by itself.

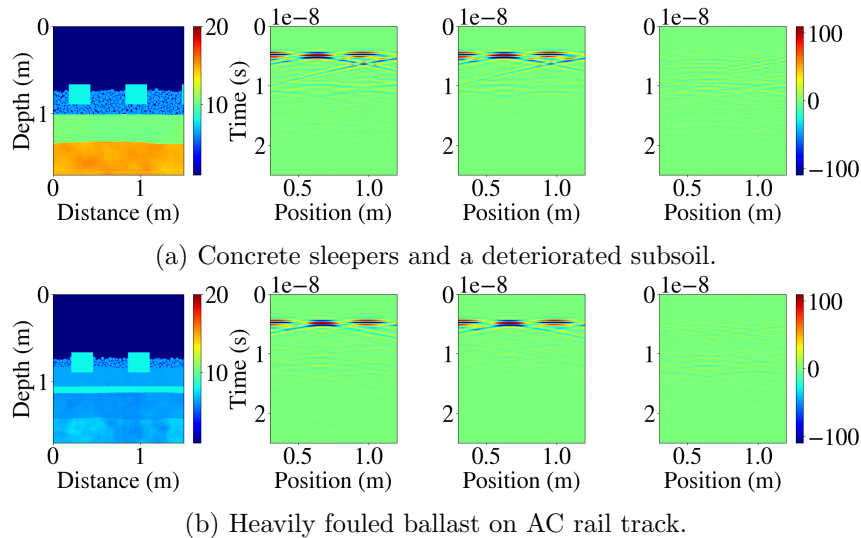


Figure 3: From left to right: sample ϵ_r geometry, ground truth, black-box model predictions and difference with ground truth.

The predictions for sleeper responses are generally accurate, with minor discrepancies and are similarly accurate across all three sleeper types. The loss values indicate that the least accurate predictions occur where groundwater is present between layers, particularly at the subgrade-subsoil interface. Samples with the highest validation loss are reported in Figure 4. In terms of computational cost at inference time, the black-box model is significantly faster than FDTD simulations: full B-scan predictions for a batch of 30 samples take approximately 300ms on our setup, compared to around 5 minutes for computing the B-scan response of a single dataset geometry on the same machine, yielding a speedup factor of approximately 30,000. This substantial speedup shifts the bottleneck in dataset generation to the creation of randomized sample geometries, resulting in an overall process speedup of around 100 times.

5 PINN MODELS

Physics-informed neural networks exploit physical PDEs and boundary conditions to apply additional loss terms to the training of NN models, with weighting coefficients λ adopted to stabilize the training process:

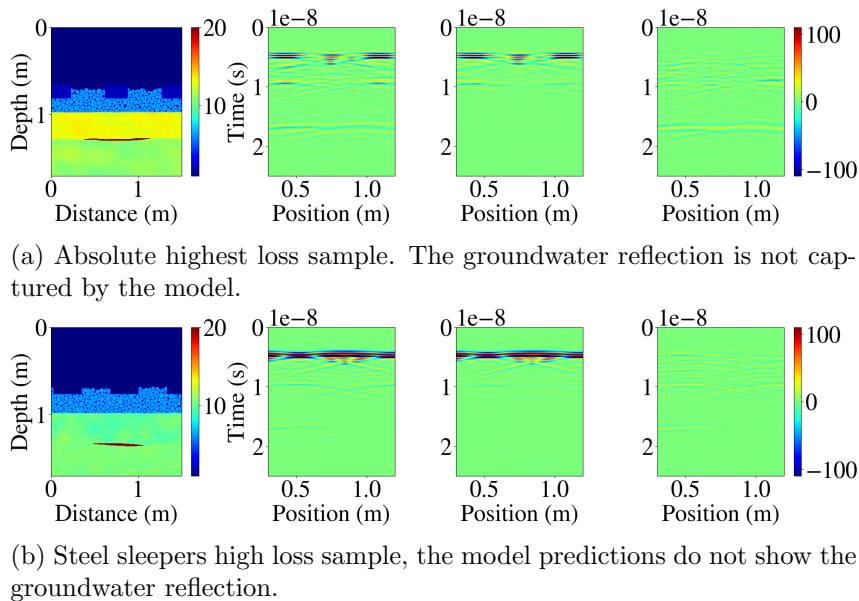


Figure 4: From left to right: sample ϵ_r geometry, ground truth, black-box model predictions and difference with ground truth of some of the samples with highest difference between the prediction and label.

$$L = \lambda_{obs}L_{obs} + \lambda_{pde}L_{pde} + \lambda_{bc}L_{bc} \quad (1)$$

where each loss component L_i is obtained by applying MSE over the residuals. We adopt the second order scalar wave equation in two dimensions for the electric field PDE [20]:

$$\frac{\partial^2 E}{\partial t^2} - \frac{1}{\mu\epsilon}\nabla^2 E + \frac{\sigma}{\epsilon}\frac{\partial E}{\partial t} = 0 \quad (2)$$

where $\nabla^2 E = \frac{\partial^2 E}{\partial x^2} + \frac{\partial^2 E}{\partial y^2}$. In training, the residuals from the PDE are minimized over an extended collocation domain, beyond that of the observations. To this end, predictions of the evolution of the electric field over space and time are produced.

5.1 MLP-based PINN architecture

We present a family of physics-informed MLP models for the prediction of electric field values in space and time, associated with increasingly complex geometry maps. All models have 3 inputs; spatial coordinates and time (x , y , t), and one output representing the electric field. The approach adopted in [19] is here reproduced for the first two experiments, namely the uniform and two-layer geometries. The derivative terms in Equation 2 are calculated using automatic differentiation. The input data was normalised into the $[0, 1]$ range and labels between $[-1, 1]$.

Uniform geometry: In the first experiment on a uniform geometry, the results shown by Zheng and Wang [19] are reproduced and confirmed by our model. The PINN model appears indeed significantly more accurate on domain extension tasks compared to its non physics-informed counterpart. As our experiment produces identical results compared with [19], these are omitted here for brevity.

Two layer geometry

We conduct a second experiment similar to the one by Zheng and Wang [19] and further confirm the results obtained on a two layer geometry. However, this scenario includes observation data that already captures the reflection of EM waves due to the layer interface. Thus, we here also investigate the case where observation data only covers the time frame before any reflection occurs with two different experiments. The first only exploits observation and physics loss, while for the second experiment, electric field boundary conditions are additionally employed at the dielectric interface. These are, in the absence of unbound surface charges [21]:

$$\hat{n} \cdot D_2 = \hat{n} \cdot D_1 \quad (a) \quad \hat{n} \times E_2 = \hat{n} \times E_1 \quad (b) \quad (3)$$

where \hat{n} represents the propagation direction of the EM waves. Since our experiments involve 2D simulation, Equation 3(a) becomes trivially satisfied, as propagation direction and D field are in each point perpendicular. Thus, only Equation 3(b) is implemented. The domain size




Figure 5: Rel. permittivity, 2-layer geometry.




Figure 6: Results on the two layer geometry at $t = 45\text{ns}$ with training snapshots between 15 and 23 ns. From left to right: ground truth, PINN predictions, NN prediction, difference between PINN prediction and ground truth.

is $20\text{m} \times 20\text{m}$ and observations are provided on the time snapshots at $t \in [15, 19, 23]\text{ns}$, as the reflection occurs approximately at $t = 25\text{ns}$. The original wave is a Ricker wavelet with a central frequency of 150MHz, emitted at the center of the simulation space. The interface between the two materials is placed horizontally 5 meters above the bottom of the domain, as shown in Figure 5. The materials have relative permittivity of 1 (above) and 5 (below). Full-batch gradient descent is used for training, coupled with an MSE loss and the Adam optimizer. Relative λ values are: $\lambda_{obs} = 1$, $\lambda_{pde} = 2 \cdot 10^{-18}$, $\lambda_{bc} = 1$. Figure 6 shows predictions of the models at $t = 45\text{ns}$. The PINN models, albeit more accurate than their non physics-informed counterparts, do not fully capture the reflected wave. Furthermore, the application of boundary conditions does not seem to substantially affect the results.

Railway track geometry

Finally, the MLP-based PINN architecture is applied to the more complex railway track geometries. The application of the model to a single sample from the A-scan dataset is shown in Figure 7. The domain size is $1.5\text{m} \times 1.7\text{m}$.

The simulation is seeded with a Ricker wavelet with central frequency 1GHz, originated 46.5cm above the center of the two concrete sleepers. Snapshots of the electric field between 2ns and 19ns were used as observations. From our experience, the training process in these conditions is unstable and no MLP model (physics-informed or not) accurately captures the full wavefield obtained from a single railway geometry. Figure 8 shows the “best” results obtained with this setup, both for PINN and non physics-informed networks. The less than satisfactory predictions may be a result of the high number of wave reflections. Zheng and Wang [19] show a similar result on their most complex geometry, where they speculate that the cause may lie in the spectral bias of NNs, which inhibits learning high frequency components of the solution in a reasonable time frame.







Figure 7: Rel. permittivity, rail track geometry.



(a) $t = 3.5\text{ns}$, notice how the PINN shows a smoother predicted wavefield w.r.t. the NN.



(b) $t = 8.5\text{ns}$, the NN learned a more precise representation of the reflected wavefield, but many reflections are already not captured by either model.



(c) $t = 14.5\text{ns}$, the PINN predicts $E \approx 0$ everywhere, while the NN partially captures the main waves at the bottom.

Figure 8: From left to right: ground truth, PINN predictions, NN predictions, PINN difference, NN difference with ground truth. Applied to the railway dataset geometry trained on snapshots between 2ns and 19ns every 1ns.

Further reasons might involve the size and architecture of the networks, which in this study is restricted to MLPs with a maximum width of 512 neurons per layer, and depth of 5 layers.

5.2 CNN-based PINN architecture

Given the complexity and image-like nature of the GPR observations, we further investigate a physics-informed CNN model, as a first attempt of this kind to the GPR simulation problem. This model takes the time coordinate as input and outputs a 2D image of the electric field at the specified time step. As no geometry data is provided, only a single dataset sample is fed to the model. A PINN model




Figure 9: Schematic representation of the CNN PINN model

schematic representation of this architecture is displayed in Figure 9. It features an initial MLP with increasing layer sizes, followed by two convolutional decoding modules.

Two notable differences are present in the computation of the PDE loss to enable efficient training of this CNN model, compared to the previous experiments involving MLPs. Namely, i) forward mode automatic differentiation is used to efficiently calculate first and second order derivatives of the E field image with respect to the input time coordinate, and ii) numerical differentiation is used on the output field to compute its second order derivatives in space.

For this experiment we used observations between 2ns and 10ns. Figure 10 shows the predictions of the two networks (physics-informed and regular) at different points in time. Both networks are able to successfully capture the training data. However, they exhibit deficiencies in sparse reconstruction and domain extension capabilities. These properties are typically expected from a PINN model and, in particular, sparse reconstruction becomes crucial for precise B-scan predictions. We speculate that the reasons for these inaccuracies lie in the numerical differentiation method used to approximate image derivatives, although further research is needed to determine the exact causes of this phenomenon.

5.3 Additional insights - 1D experiment

To further investigate the reasons behind the downfall of the physics-informed learning, when combined with the CNN architecture, an additional experiment is designed. A 1D wave propagates towards the positive x direction in time. Three different models are trained on this data: an MLP with x and t as input, a CNN model adapted to predict 1D sequences, and a further MLP that takes the time coordinate as input and outputs a vector predicting the full wave at the corresponding time, which is denoted *discrete MLP*. Five observation time steps were used for training, corresponding to $t = [0, 1, 2, 3, 4]s$, and wave speed of 16 m/s. The spatial domain is 512m for the MLP model, with 20000 (x, t) collocation points spread in $t \in [0, 40]s$, and corresponding 512 pixels for both the CNN and discrete MLP architectures, with 2000 collocations in the same time domain. Figure 11 reveals that the MLP models show remarkable sparse interpolation capabilities, while both the CNN and the discrete MLP models are less accurate. For domain extension, the MLP predictions are again perfect, while both discrete MLP and CNN models are unable to capture the wave evolution. The results hint that the inaccuracy produced by the physics-informed CNN may stem from the discrete nature of the learning task of the CNN, which does not work well with the physics-informed training.

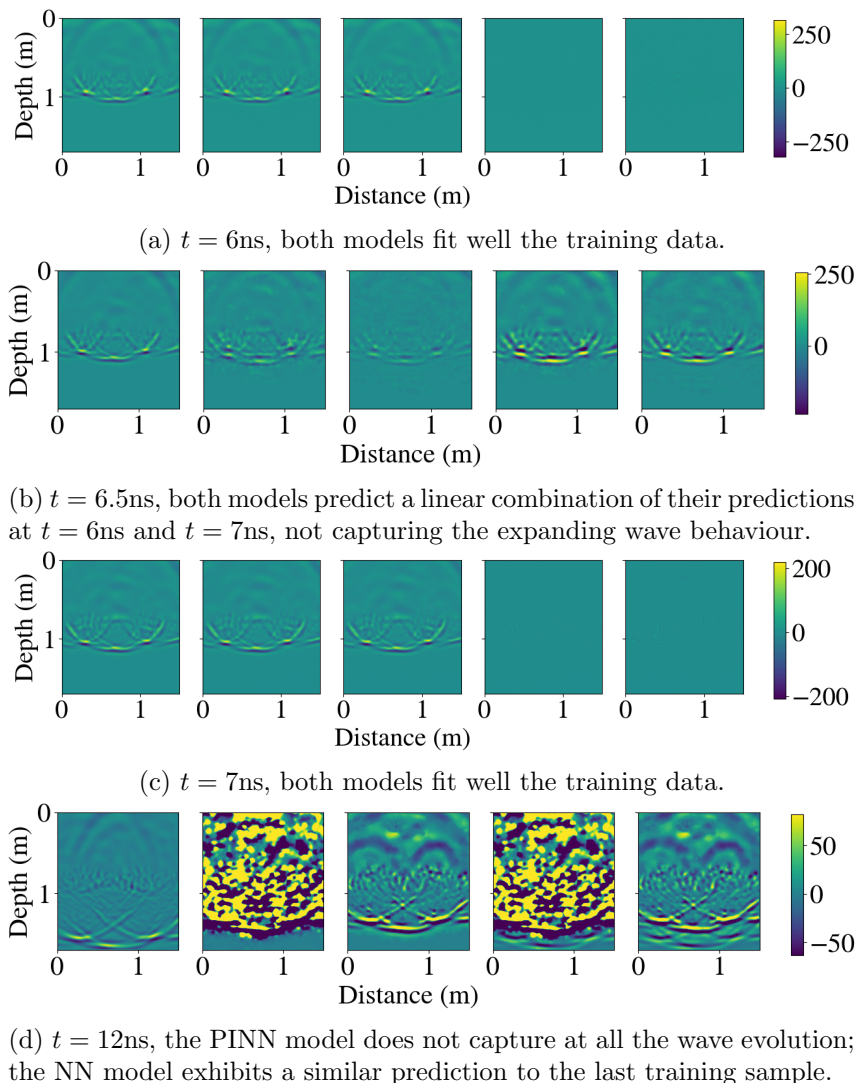


Figure 10: From left to right: ground truth, PINN and NN predictions, PINN and NN difference to ground truth. Applied to the CNN-based model on the railway dataset geometry.

6 DISCUSSION AND CONCLUSIONS

This study illustrates the generation of two datasets (including A-scan and B-scan outputs) of diverse and realistic GPR railway tracks using the open-source library `gprMax`, based on information provided by the Swiss Federal Railways. These datasets, while realistic, have limitations, such as relying on a 2D representation of railway track geometry, excluding rails and antenna models. The B-scan dataset is used to train a black-box encoder-decoder CNN, yielding promising results with predictions that almost always closely match FDTD simulations. The model computational cost at inference time is approximately 30000 times lower than `gprMax`, achieving an overall speedup of about 100 times in dataset generation. In addition, we show the application of the A-scan dataset to MLP- and CNN- based PINN architectures, showing their limitations, in particular in terms of interpolation and extrapolation from the training domain.




Figure 11: Results from the 1D wavefield experiment on different PINN models. The red dotted line indicates the last training observation at $t = 4$ seconds.

Future work will focus on applying the trained black-box model to generate large-scale datasets, facilitating subsequent monitoring methods for railway infrastructure maintenance. Enhancing dataset realism to account for external objects and a comprehensive 3D representation, including antenna, rail, and train carriage reflections, is crucial. Research on PINNs should explore the impact of boundary conditions at layer interfaces, possibly influenced by our 2D representation. Further research is needed on applying PINNs to complex geometries, particularly for modeling electromagnetic waves. Moreover, the hereby presented PINN models with discretized output domains demonstrate limited success, suggesting a need to explore the causes of this issue. An extension of the PINN framework to CNN architectures could be a promising such direction.

7 DATA AVAILABILITY AND ACKNOWLEDGMENTS

The code used to generate the datasets, train and evaluate the models in this work is publicly available online at <https://github.com/ThomasRigoni7/PINN4GPR>. The generated datasets and pre-trained ML models are available on request. The authors acknowledge the support of the Swiss Federal Railways (SBB) as part of the ETH Mobility Initiative project REASSESS.

References

- [1] C. Hoelzl, V. Dertimanis, M. Landgraf, L. Ancu, M. Zurkirchen, and E. Chatzi, “On-board monitoring for smart assessment of railway infrastructure: A systematic review,” in *The Rise of Smart Cities*, 2022, pp. 223–259.
- [2] A. M. Zaremski, G. T. Grissom, and T. L. Euston, “On the use of ballast inspection technology for the management of track substructure,” *Transportation Infrastructure Geotechnology*, vol. 1, no. 1, pp. 83–109, Mar. 2014.
- [3] S. Wang, G. Liu, G. Jing, Q. Feng, H. Liu, and Y. Guo, “State-of-the-art review of ground penetrating radar (GPR) applications for railway ballast inspection,” *Sensors*, vol. 22, no. 7, 2022.
- [4] G. Arcieri, T. Rigoni, C. Hoelzl, D. Haener, and E. Chatzi, “Ground penetrating radar for moisture assessment in railway tracks: An experimental investigation,” in *Proceedings of the 10th European Workshop on Structural Health Monitoring*, 2024.
- [5] C. Charoenwong, D. Connolly, P. Alves Costa, P. Galvín, and A. Romero, “The effect of ballast moisture content and fouling index on railway track settlement,” *Transportation Geotechnics*, vol. 45, p. 101 193, 2024.

- [6] L. Wang, M. Meguid, and H. S. Mitri, “Impact of ballast fouling on the mechanical properties of railway ballast: Insights from discrete element analysis,” *Processes*, vol. 9, no. 8, p. 1331, 2021.
- [7] N. S. Kahil *et al.*, “Automatic analysis of railway ground penetrating radar: Using signal processing and machine learning approaches to assess railroad track substructure,” *Transportation Research Procedia*, vol. 72, pp. 3008–3015, 2023.
- [8] C. Warren, A. Giannopoulos, and I. Giannakis, “Gprmax: Open source software to simulate electromagnetic wave propagation for ground penetrating radar,” *Computer Physics Communications*, vol. 209, pp. 163–170, 2016. [Online]. Available: <https://www.gprmax.com/>.
- [9] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational physics*, vol. 378, pp. 686–707, 2019.
- [10] M. Haywood-Alexander, W. Liu, K. Bacsa, Z. Lai, and E. Chatzi, “Discussing the spectra of physics-enhanced machine learning via a survey on structural mechanics applications,” *arXiv preprint arXiv:2310.20425*, 2023.
- [11] B. Moseley, A. Markham, and T. Nissen-Meyer, *Solving the wave equation with physics-informed deep learning*, 2020. arXiv: 2006.11894 [physics.comp-ph].
- [12] M. Rasht-Behesht, C. Huber, K. Shukla, and G. E. Karniadakis, “Physics-informed neural networks (PINNs) for wave propagation and full waveform inversions,” *Journal of Geophysical Research: Solid Earth*, vol. 127, no. 5, Apr. 2022.
- [13] P. Koyan and J. Tronicke, “3D modeling of ground-penetrating radar data across a realistic sedimentary model,” *Computers & Geosciences*, vol. 137, p. 104422, 2020.
- [14] 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*, vol. 9, no. 1, pp. 37–51, 2016.
- [15] I. Giannakis, A. Giannopoulos, C. Warren, and N. Davidson, “Numerical modeling and neural networks for landmine detection using ground penetrating radar,” Jul. 2015.
- [16] I. Giannakis, A. Giannopoulos, and C. Warren, “A machine learning approach for simulating ground penetrating radar,” in *2018 17th International Conference on Ground Penetrating Radar (GPR)*, 2018, pp. 1–4.
- [17] I. Giannakis, A. Giannopoulos, and C. Warren, “A machine learning-based fast-forward solver for ground penetrating radar with application to full-waveform inversion,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 7, pp. 4417–4426, 2019.
- [18] Q. Dai *et al.*, “A deep learning-based gpr forward solver for predicting b-scans of subsurface objects,” *IEEE Geoscience and Remote Sensing Letters*, vol. 19, pp. 1–5, 2022.
- [19] Y. Zheng and Y. Wang, “Ground-penetrating radar wavefield simulation via physics-informed neural network solver,” *Geophysics*, vol. 88, no. 2, KS47–KS57, Feb. 2023.
- [20] S. Zarei, B. Oskooi, N. Amini, and A. R. Dalkhani, “2D spectral element modeling of GPR wave propagation in inhomogeneous media,” *Journal of Applied Geophysics*, vol. 133, pp. 92–97, 2016.
- [21] ETH Zürich, *Material boundaries*. [Online]. Available: <https://ethz.ch/content/dam/ethz/special-interest/itet/photonics-dam/documents/lectures/EandM/MaterialBoundaries.pdf>.