geoconvention
Calgary • Canada • May 12-14 2025

Towards applying physics informed neural networks in realistic mining scenarios

Alejandro Jaimes, Miro Doring, Dmitriy Malovichko, and Alex Rigby
Institute of Mine Seismology

Summary

Travel-time tomography is a non-invasive technique that uses the arrival time of seismic waves to infer the velocity structure of the medium that these waves traverse. While commonly used for exploration purposes, tomography has the potential to estimate temporal changes in velocity which can then be associated to changes in the stress state of the rockmass. Traditionally, travel-time tomography has been performed using linear, pseudo-linear, or adjoint-state-based methods. More recently, physics informed neural networks (PINNs) have emerged as an alternative to classical tomography methods due to their speed and accuracy, particularly in large-scale three-dimensional settings. In this paper, we showcase the use of PINNs for a realistic (yet synthetic) mining scenario. We illustrate improvements to standard PINNs-based tomography by setting bounds on allowable slowness and travel-times (and their gradients), as well as non-uniform sampling which accelerates convergence and accuracy of the inversion. We find PINNs to outperform classical methods when dealing with a large number of triggers (>100k). Further, PINNs easily allow incorporating geological, geotechnical and geophysical priors (such as known lithologies, shapes or boundaries) which are of importance for accurately interpreting velocity changes.

Theory & Methods

The eikonal equation for an isotropic medium, following the notation of (Cerveny,2001), can be written as

$$\left(\left|\nabla T\left(x_s, x\right)\right|\right)^2 = s^2\left(x\right),$$

where T is the travel time field, and s is the slowness field. The vectors x_s and x denote the three-dimensional position vectors for the source and an arbitrary location within the medium of interest, respectively.

The eikonal equation is well known to suffer of a source singularity. To deal with this problem, the travel time field is typically rewritten in an additive (Taufik et. al., 2023) or multiplicative form (Treister and Haber, 2016). Here we focus on the latter form, by which the travel-time field can be split as

$$T\left(x_s, x\right) = R\,\tau\left(x_s, x\right),$$

where $R = \left|x_s - x\right|$, and τ is a scalar field relating the distance R to the actual travel-time field T.

Inserting this expression for the travel-time field into the isotropic eikonal equation above gives the modified isotropic eikonal equation

$$\left|R\nabla\tau\left(x_s,x\right)+\tau\left(x_s,x\right)\nabla R\right|^2=s^2\left(x\right).$$

In the context of travel-time tomography, we aim to use either the modified or original eikonal equation to determine the slowness field $s\left(x\right)$ from the observed travel time field $\tau\left(x_s,x\right)$ (or $T\left(x_s,,x\right)$. Notice, however, that we only have access to the true travel time field at a collection of discrete receivers, meaning that the accuracy of the retrieved slowness model is limited by the ray coverage between source and receiver locations.

Following the seminal work by Raissi et. al. (2019), PINNs have emerged as a powerful tool for solving forward and inverse problems involving nonlinear partial differential equations. In this paper we focus only on the inverse formulation. To set-up a PINNs-based inverse problem for the isotropic eikonal equation we first reformulate the modified eikonal equation as

$$\left|R\nabla N_\tau\left(x_s,x\right)+N_\tau\left(x_s,x\right)\nabla R\right|^2=N_s^2\left(x\right),$$

where we have substituted the travel-time (τ) and slowness (s) scalar fields, by a travel-time (N_τ) and slowness (N_s) neural network approximation, respectively. The travel-time and slowness neural networks take the vector $\left(x_s,x\right)$ and $\left(x\right)$ as their inputs, respectively. To facilitate the inversion we rewrite the output of the neural networks as follows

$$N_s\left(x\right)=\sigma\left(\widehat{N}_s\left(x\right)\right)*\left(s_{max}-s_{min}\right)+s_{min},$$

$$N_\tau\left(x_s,x\right)=\sigma\left(\widehat{N}_\tau\left(x_s,x\right)\right)*\left(s_{max}-s_{min}\right)+s_{min},$$

where s_{min} and s_{max} are minimum and maximum slowness values, and σ is the sigmoid function whose range is $\left(0,1\right)$. The output of N_s is between given s_{min} and s_{max} values. Similarly, the output of N_τ is between s_{min} and s_{max}. This means that the predicted travel-time T is constrained to be between $R\,s_{min}$ and $R\,s_{max}$. Notice that the constraints on the bounds are motivated by the work of (Grubas et. al., 2023) who first proposed imposing bounds on travel-times for solving the eikonal equation. Such improvements by Grubas et. al. (2023) can be similarly applied to the inverse problem of determining both the travel-time and slowness fields from sparse travel-time measurements, and do in fact improve convergence and accuracy. Since $s\left(x\right)$ is bounded in the range $\left[s_{min},s_{max}\right]$ it makes sense for $\left|\nabla T\right|$ to also be bounded on the same range. Bounding the travel-time network on said range does not suffice to bound its gradient. To speed up convergence and improve accuracy, we apply a clipping operator $\eta\left(\alpha,s_{min},s_{max}\right)$ to the modified eikonal equation, which equals α for $\alpha\in\left[s_{min},s_{max}\right]$ and defaults to the respective bounds, otherwise.

Thus, we aim to solve

$$\eta\left(\left|R\nabla N_\tau\left(x_s,x\right)+N_\tau\left(x_s,x\right)\nabla R\right|^2,s_{min},s_{max}\right)=N_s^2\left(x\right),$$

The goal with PINNs-based eikonal tomography is then to find the neural networks $\widehat{N}_\tau$ and $\widehat{N}_s$ which most closely satisfy the clipped eikonal equation, subject to RN_τ matching the recorded travel-times at discrete receivers, that is $RN_\tau\left(x_s,x_r\right)\approx T^d\left(x_s,x_r\right)$ where T^d is the first arrival travel-time recorded at a receiver at location x_r due to a source at location x_s. To this end, we construct the following objective function

$$J\left(\theta\right)=E_{x'\sim\rho}E_{x\sim\beta}\left(\eta\left(\left|R\nabla N_\tau\left(x_s,x;\theta\right)+N_\tau\left(x_s,x;\theta\right)\nabla R\right|^2,s_{min},s_{max}\right)^2-N_s^2\left(x;\theta\right)\right)^2$$
$$+\frac{1}{N_{src}}\frac{1}{N_{rec}}\sum_{i=1}^{N_{rec}}\sum_{j=1}^{N_{src}}\left(RN_\tau\left(x_s,x_r;\theta\right)-T^d\left(x_j,x_i\right)\right)^2$$

which we aim to minimise. The integers i,j denote the receiver and source indices, respectively. The parameter $\theta=\left(\theta_\tau,\theta_s\right)$ denotes the set of weights for the travel-time and slowness neural networks, respectively, that we aim to optimise over. More specifically, we want to find

$$\theta^{opt}=argmin_{\theta\in\Theta}J\left(\theta\right).$$

Thus, we seek weights that minimise both the travel-time data misfit as well as the clipped eikonal equation residual. In this paper, we use ADAM, a popular and efficient method for effective stochastic optimisation. Once the optimal weights θ^{opt} have been determined, the optimal slowness model is found by evaluating the slowness neural network at said weights.

Results

Due to limitations regarding permission of field data, in this abstract we only show a simple checkerboard test and a synthetic test with a toy velocity model. The checkerboard model is constructed using an average velocity of $v_p=4.31$ km/s with alternating cells of lower and higher velocities. Checkerboard tests are useful to test how tomographic algorithms may be able to retrieve variations of seismic velocities at nearby locations. The figure below shows a comparison of the checkerboard model (top left) and retrieved models using classical linear tomography (top right), the simultaneous iterative reconstruction technique (bottom left), and PINNs (bottom right). The location of sources (not shown here) are synthetically simulated based on the evolution of stress acrosss the mining front (in blue). The receivers (not shown here) are placed accordingly to emulate a realistic mining scenario. Thus, the ray density is not uniform in the domain of interest. Clearly PINNs outperforms the other two tomography methods not only in the retrieved values but also the shape of the anomalies.

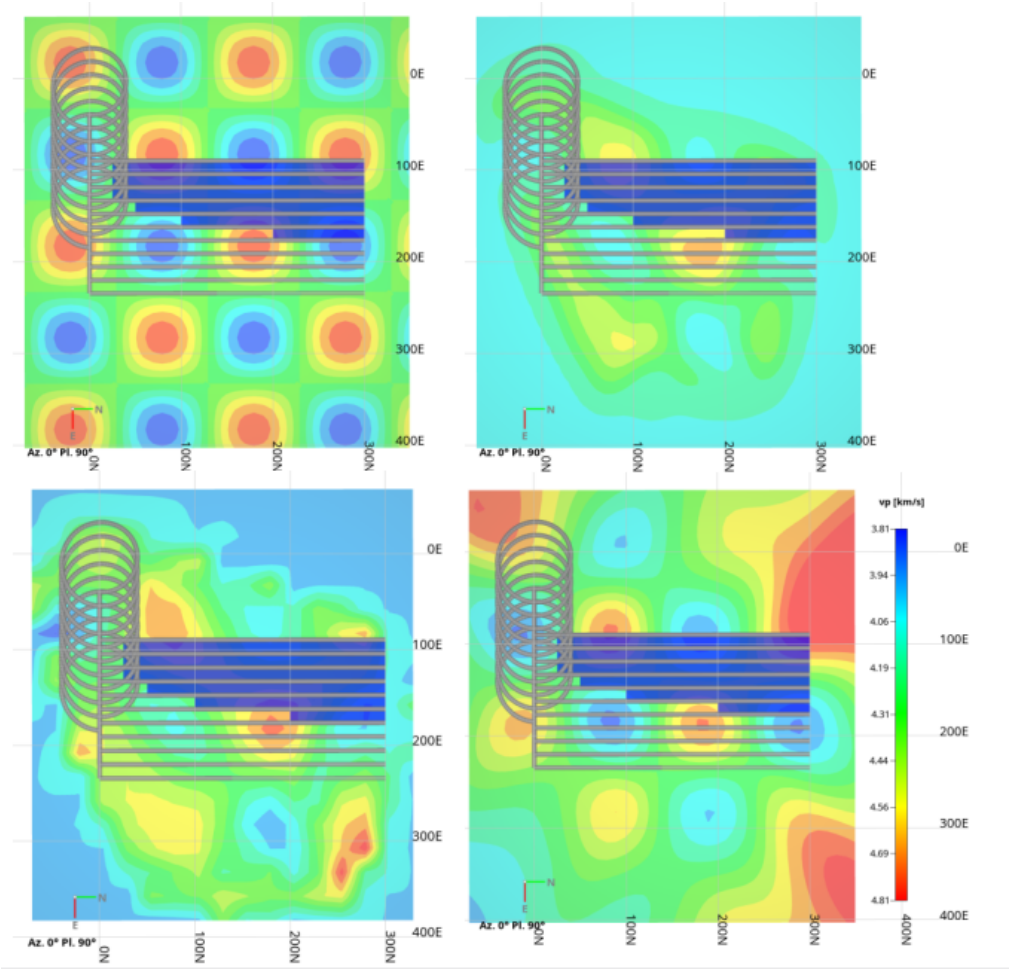

Figure 1: Comparison of exact checkerboard pattern (top left) with results from linear tomography (top right), SIRT method (bottom left), and PINNs (bottom right).

As a simple test case we consider a velocity model derived from a linear relationship with maximum compressive stress σ_1, which is obtained from stress modelling of the same mining sequence shown in the figure above. This scenario is of importance in the mining industry as seismic hazard can developed in areas of significant stress concentration. The figure below shows the retrieved velocity model. Inside the mining stopes the velocity is low (and so is the stress). In the perimeter of the mining front the velocity rapidly increases which agrees with stress concentration from numerical modelling. Overall, the algorithm provides a realistic and interpretable image which may be of further use for geotechnical analysis

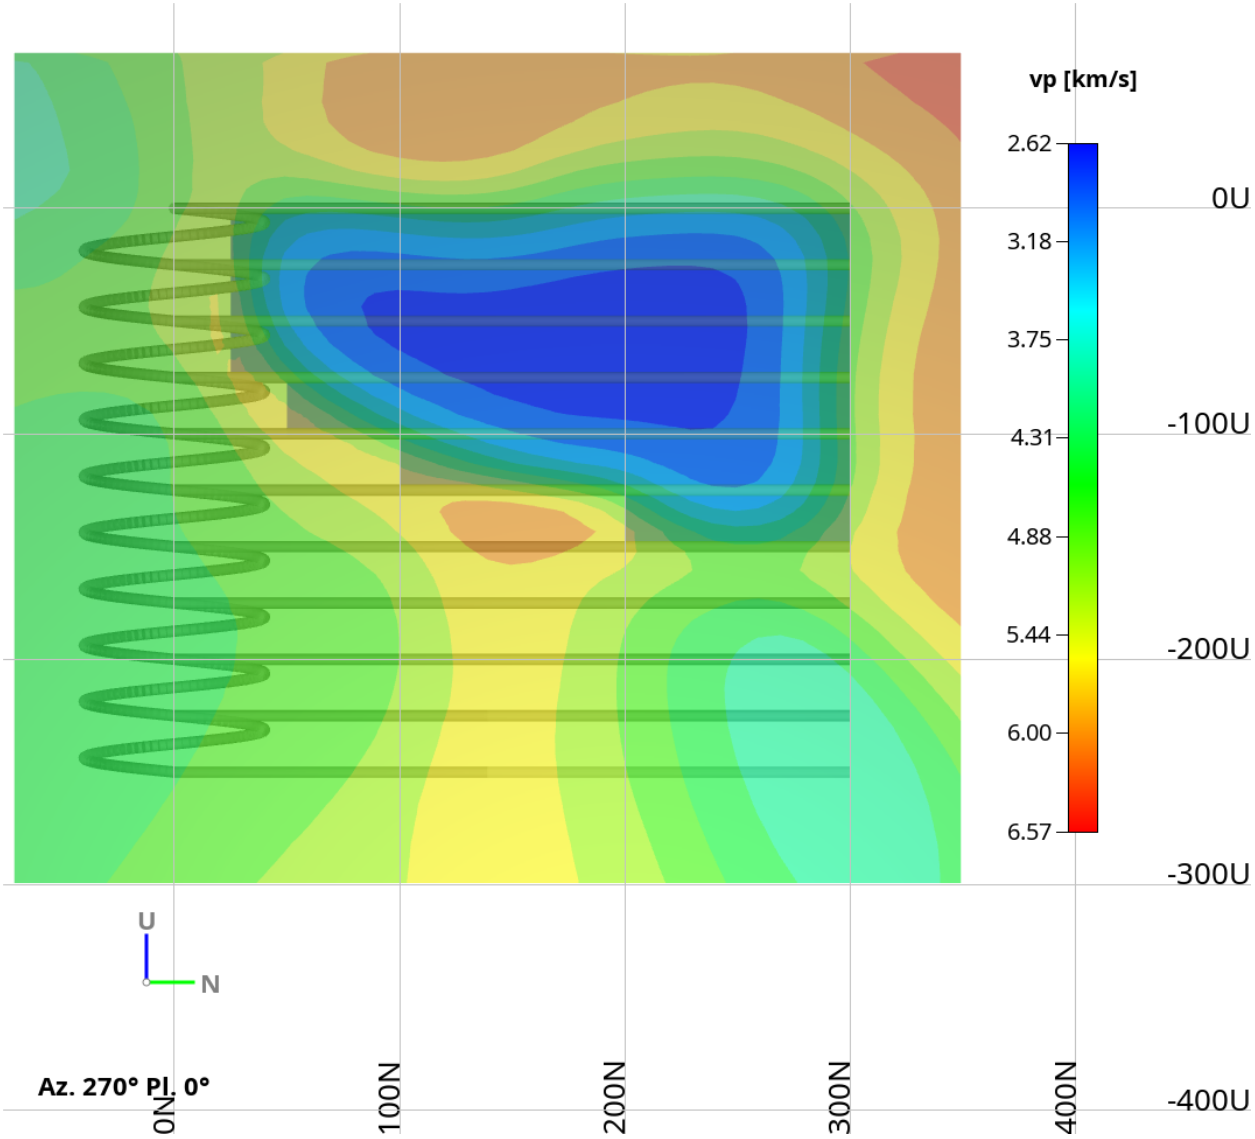

Figure 2: West view of retrieved velocity model using PINNs.

Conclusions

In the context of travel-time tomography PINNs have emerged as a possible alternative to classical methods. PINNs are natural for the travel-time tomography problem as the neural network can be used to learn the forward model (the solution of the eikonal equation), and inverse problem (retrieving the velocity model of interest) simultaneously in the learning process. Indeed, this simply requires constructing an objective function which consists of a data misfit term penalizing deviations from observations as well as the PDE residual penalizing the deviations from the eikonal equation. PINNs-based tomography is accurate and relatively efficient for large data sets, as compared to classical methods. This is because one does not need to store model derivatives but instead neural network weights. Numerical tests with field data (not shown here) indicate that for a data set of 100k source-receiver pairs the tomographic inversion takes less than an hour on a personal computer with an NVIDIA RTX 3500 ADA graphics card. To enhance the convergence an accuracy of the inversion one can impose bounds on the retrieves slowness and travel-times (and their gradients). Naturally, this improves the retireved models as the model space is significantly shrunk. Further, the sampling algorithm that is used in PINNs has a significant effect when dealing with sparse geometries and should be taken under consideration when buidling PINNs-based tomography. In the next section we provide details about the importance of sampling and advocate for physically-informed sampling schemes. While not mentioned explicitly here, because PINNs are based on a cloud point representation of the domain of interest, rather than a grid (as in classical methods), it is straight-forward to incorporate geological, geotechnical and geophysical priors (such as known lithologies, shapes or boundaries). Such priors are of importance for accurately representing velocity changes.

Novel/Additive Information

The algorithm presented herein has potential to be useful in realistic scenarions (not shown here due to issues with permission). Incorporation of bounds on the predicted slownesses and travel-times (and their gradients) significantly speeds up the convergence of PINNs and enhances their accuracy, which is of importance if PINNs are to be used for tomography in realistic settings. Furthermore, while not discussed in detail here, the sampling of the collocation points that are used to enforce the eikonal equation is not arbitrary. While uniform sampling is the most optinal approach when the ray density is pseudo-uniform, it is not optimal when dealing with very sparse receiver geometries, as is the case in underground mining. For such geometries, we suggest instead to use a physically-informed sampling scheme based on Fresnel volumes computed in an arbitrary background medium. In this way, one can avoid excessive sampling in areas that are not sampled by seismic rays, while focusing on areas that seismic rays do indeed traverse. The figure below shows a comparison of results obtained with PINNs (uniform versus Fresnel-based sampling) for the same checkerboard model as in the previous section. Notice that Fresnel-based sampling outperforms uniform sampling since we are focusing on areas of the model which are actually sampled by seismic rays, rather than excessively training the networks on areas where there is no ray coverage.

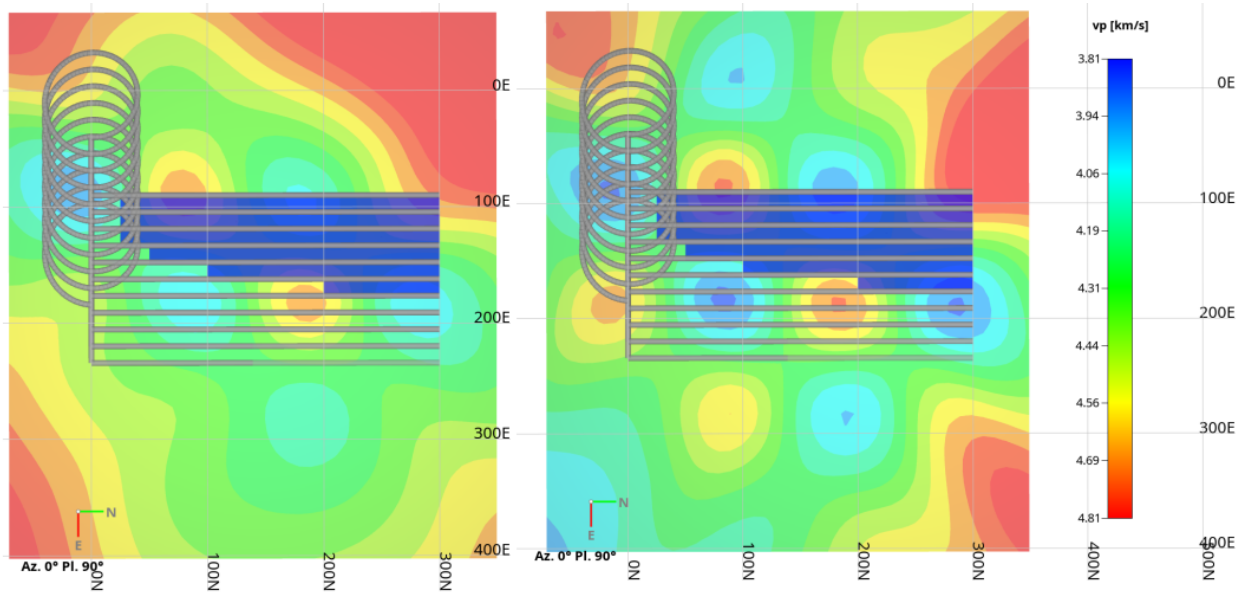

Figure 3: Comparison of inverted checkerboard model for uniform (left) and Fresnel (right) sampling scheme.

Acknowledgements

We thank colleagues at the Institute of Mine Seismology for meaningful discussions.

References

1. Cerveny. Seismic Ray Theory. Cambridge University Press, 2001.

2. M. Taufik, T. Alkhalifah, and U. Waheed. A stable neural network-based eikonal tomographyusing hard-constrained measurements. Authorea Preprints, 2023.

3. E. Treister and E. Haber. A fast marching algorithm for the factored eikonal equation. Journal of Computational Physics, 324:210–225, 2016.

4. 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, 378:686–707, 2019.

5. Serafim Grubas, Anton Duchkov, and Georgy Loginov. Neural eikonal solver: Improving accuracy of physics-informed neural networks for solving eikonal equation in case of caustics. Journal of Computational Physics, 474:111789, 2023.