Modeling Cesium-137 Transport in Soils Using Physics-Informed Neural Networks

Ilona Kopintsu
ISP RAS
Moscow, Russia
e-mail: redlonelymountain@gmail.com

Sergei Strijhak
ISP RAS
Moscow, Russia
e-mail: s.strijhak@ispras.ru

***Abstract*—Accurate modeling of radionuclide migration in soils is essential for environmental safety and post-accident remediation. Cesium-137 (^{137}Cs) is of particular concern due to its ∼30-year half-life and strong sorption to soil constituents. We develop and assess a hybrid framework that couples numerical moisture dynamics with Physics-Informed Neural Networks (PINNs) to simulate ^{137}Cs transport in variably saturated soils. Water content $\theta(z,t)$ is obtained from Richards-type simulations, whereas contaminant transport is governed by an advection–diffusion–reaction model with kinetic sorption and decay. The PINN approximates $C(z,t)$ from inputs (z,t,θ,soil); derivatives are computed by automatic differentiation and used to form a PDE residual. Three soil-type encodings are compared: scalar, one-hot, and trainable embeddings. Performance is evaluated with structural similarity (SSIM), root-mean-square error (RMSE), and mean absolute error (MAE), with a focus on generalization to unseen soils. Explicit soil encoding substantially improves accuracy and physical consistency. One-hot encoding achieves the lowest residuals on seen soils, while embeddings provide the best zero-shot transfer to unseen soils. Residual maps confirm improved PDE adherence when soil-aware inputs are included. The encoding strategy is a key design choice for physics-informed modeling of heterogeneous porous media. The proposed framework provides a foundation for soil-aware PINNs and can be extended to dual-network setups that jointly predict $\theta(z,t)$ and $C(z,t)$.**

***Keywords*—PINN, ^{137}Cs, sorption, unsaturated flow, advection–diffusion–reaction, embeddings, residual maps.**

I. Introduction

Radionuclide migration in soil is driven by the interplay of unsaturated flow, sorption–desorption kinetics, and radioactive decay. Classical numerical solvers of advection–dispersion–reaction (ADR) equations require detailed parameterization (retention curves, conductivity, kinetic rates), which is challenging in heterogeneous soils and costly for long horizons [1].

Physics-Informed Neural Networks (PINNs) embed governing PDEs into the learning objective, and can leverage sparse data while keeping physical consistency [2]. Prior works focused mainly on water flow and infiltration; there is less evidence on contaminant transport with sorption and decay under heterogeneous soils [3], [4].

This paper couples numerical $\theta(z,t)$ with a PINN that predicts liquid-phase $C(z,t)$ conditioned on soil type. We compare soil-encoding strategies (none, one-hot, embeddings) and evaluate accuracy, physical fidelity via PDE residuals, and generalization to unseen soils.

II. Methodology

A. Governing Equations

The modeling framework is hybrid: the Richards equation is first solved numerically to obtain the transient water content $\theta(z,t)$ and pore-water velocity $v(z,t)$. These fields are then used as dynamic inputs to the contaminant transport model.

Moisture dynamics follow the Richards equation:

$$\frac{\partial\theta}{\partial t} = \frac{\partial}{\partial z}\left[D(\theta)\frac{\partial\theta}{\partial z}\right] - \frac{\partial K(\theta)}{\partial z}, \tag{1}$$

where $D(\theta)$ is the soil water diffusivity and $K(\theta)$ the hydraulic conductivity, parameterized by the van Genuchten–Mualem model.

The transport of dissolved cesium-137 is then described by an advection–diffusion–reaction (ADR) equation with additional sorption and decay terms:

$$\theta\,\frac{\partial C}{\partial t} = D_C\frac{\partial^2 C}{\partial z^2} - v\,\frac{\partial C}{\partial z} - \lambda C + \frac{1}{\theta}k_f(q_f - \theta C) + k_s(q_s - \theta C), \tag{2}$$

with two-site sorption kinetics:

$$\frac{dq_f}{dt} = k_f(\theta C - q_f), \qquad \frac{dq_s}{dt} = k_s(\theta C - q_s). \tag{3}$$

Here, λ is the radioactive decay constant. Boundary conditions are set as $C(0,t) = C_0$ at the surface and $\frac{\partial C}{\partial z}\big|_{z=L} = 0$ at the bottom. This coupling ensures that contaminant migration is consistently driven by the evolving water content and flow fields [5].

B. Soil II (Black Soil): Physical Setup and Parameters

Soil II represents a medium-textured "black soil" with balanced infiltration and retention. Moisture fields $\theta(z,t)$ and pore-water velocity $v(z,t)$ are obtained from the Richards simulation on a 1D column with depth $L = 100\,\text{cm}$ over a horizon $T = 14400\,\text{min}$. These fields drive the ADR transport in (2). Boundary conditions are $C(0,t) = C_0$ (surface source) and $\partial C/\partial z|_{z=L} = 0$ (no-flux at depth).
The hydrodynamic dispersion and advection entering (2) are consistently derived from the Richards solution; no additional calibration to $\theta(z,t)$ is performed.

https://doi.org/10.51408/csit2025_62

C. PINN Formulation

Inputs are $(z, t, \theta, \text{soil})$. We compare:

1) **Baseline:** no soil encoding,
2) **One-hot:** categorical soil representation,
3) **Embeddings:** trainable continuous soil vectors.

The loss is

$$\mathcal{L} = \mathcal{L}_{\text{data}} + \lambda_{\text{phys}} \mathcal{L}_{\text{PDE}}, \tag{4}$$

The loss function combines two components: a data mismatch term, defined as the mean squared error between predicted and reference concentrations, and a physics-based term, defined by the residual of the ADR equation (2). Training is performed with the Adam optimizer, using mini-batches that sample across space, time, and soil types to ensure balanced coverage of the domain.

D. Numerical Setup and Training Details

Richards simulations (1D, $z \in [0, L]$) provide $\theta(z, t)$ on a uniform grid; ADR is solved explicitly with central diffusion and upwind advection to generate reference $C(z, t)$. PINNs are trained for 3000 epochs (lr 10^{-3}); Batch size $\approx$ 5k points per soil per step. Baseline/embedding models use 3×128 Tanh; one-hot uses 4×64 ReLU. First experiment predicts $\log C$ (stability/positivity); later variants predict normalized C. SSIM is the primary metric (structure), complemented by RMSE and MAE.

III. Experimental Results for Soil II

We evaluated the three PINN configurations (baseline, one-hot, embedding) on Soil II, which represents a medium-textured black soil with balanced infiltration and retention dynamics.

A. Baseline Model (No Soil Encoding)

The baseline PINN, trained without explicit soil input, produced overly smoothed Cs-137 distributions and failed to capture the depth-dependent retention dynamics. Breakthrough curves were delayed and the adsorption phases underrepresented, reflecting the model's inability to distinguish soil-specific features.

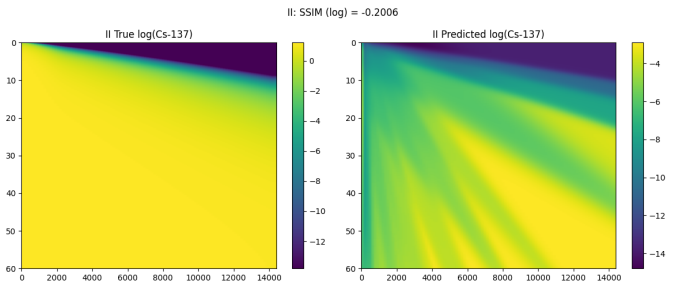

Fig. 1: Baseline PINN prediction for Soil II

B. One-Hot Encoding

With one-hot encoding, the PINN sharply captured migration fronts and retention delays. Residual maps showed uniformly low values, indicating strong agreement with the governing PDE. This configuration achieved the best accuracy for Soil II among the tested models.

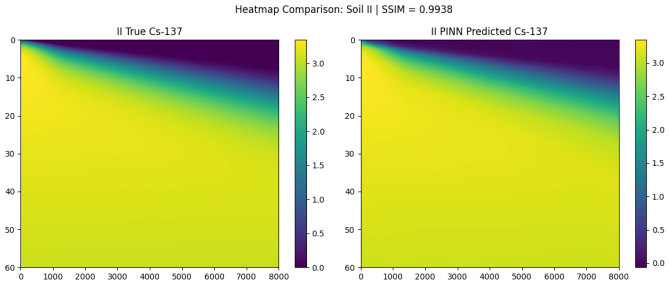

Fig. 2: One-hot encoded PINN prediction for Soil II

C. Embedding-Based Encoding

The embedding PINN generalized soil behavior through trainable continuous vectors. For Soil II, it reproduced overall plume dynamics and concentration gradients with slightly reduced accuracy compared to the one-hot model. However, embeddings provide smoother latent representations, making them suitable for extrapolation to unseen soils.

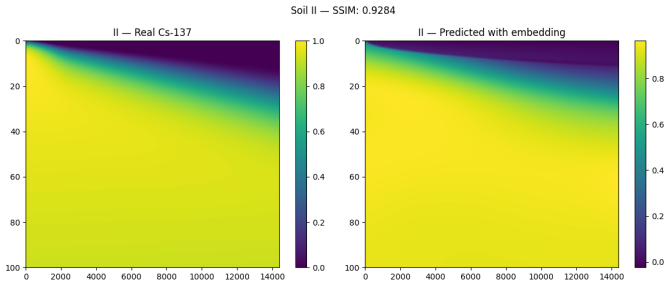

Fig. 3: Embedding-based PINN prediction for Soil II

D. Physical Consistency (Residual Maps)

Residual maps are computed by substituting $\hat{C}$ into (2). For the one-hot model (Fig. 4), residuals remain low and spatially uniform; the histogram is sharply peaked near zero, indicating good compliance with governing dynamics.

E. Temporal Profiles

Depth profiles at fixed times quantify local accuracy. Fig. 5 shows Soil II at t=5000 min; the one-hot model tracks the migration front and near-surface plateau with low bias, consistent with Table I.

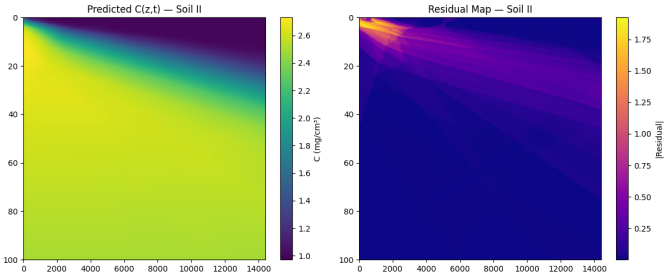

(a) Predicted vs. residual

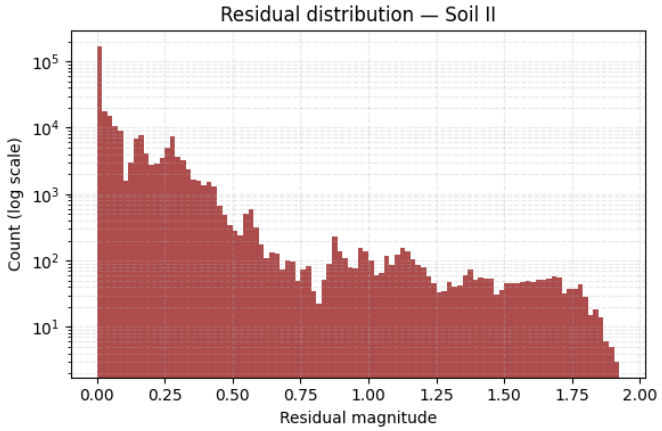

(b) Histogram

Fig. 4: PDE residual analysis for one-hot PINN on Soil II

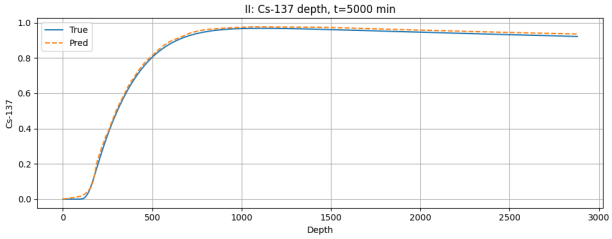

Fig. 5: Depth profile at $t = 5000$ min for Soil II (reference vs. PINN)

On Soil II, the one-hot encoded PINN achieved the highest fidelity, while embeddings offered competitive results with the added advantage of generalization. The baseline model underperformed, confirming the necessity of soil-aware inputs.

F. Training Dynamics (Soil II)

Figure 6 compares training curves for Soil II across the three PINN configurations. The baseline (no soil input) converges slowly and plateaus at a higher loss. The one-hot model achieves the fastest and lowest convergence, while the embedding model reaches an intermediate final loss, reflecting a small accuracy trade-off for improved transfer to unseen soils observed elsewhere.

(a) Baseline

(b) One-hot

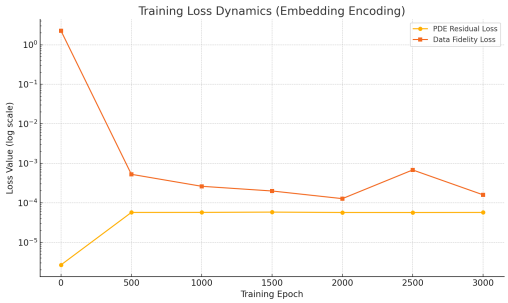

(c) Embedding

Fig. 6: Training loss on Soil II

TABLE I: Error metrics (RMSE and MAE) by soil type and time slice

Soil	Time (min)	RMSE	MAE
I	1000	1.7757e-02	4.1091e-03
I	5000	9.6885e-03	5.2479e-03
I	10000	8.3658e-03	5.7258e-03
II	1000	8.7044e-03	3.2473e-03
II	5000	6.0035e-03	3.2291e-03
II	10000	4.5601e-03	3.0623e-03
III	1000	1.7153e-02	4.8378e-03
III	5000	6.7036e-03	4.0735e-03
III	10000	5.5133e-03	4.5499e-03

Summary of Findings

The study demonstrates that the performance of Physics-Informed Neural Networks (PINNs) in modeling ^{137}Cs transport is strongly dependent on how the soil type is represented in the input. The baseline configuration, which omits explicit soil encoding, systematically underestimates retention effects and produces overly smoothed concentration fronts. This limitation reflects the model's inability to capture the soil-

specific variability of sorption–desorption kinetics and depth-dependent migration.

The one-hot encoded model substantially mitigates these issues, enabling the network to reproduce adsorption fronts, retention delays, and breakthrough curves with much greater fidelity. By conditioning predictions on categorical soil types, the model achieves the lowest residuals among the tested variants. However, this discrete representation restricts interpolation between soils and limits its ability to generalize beyond those observed during training.

The embedding-based model provides a more continuous and flexible representation. By mapping soil categories into a latent vector space, embeddings allow the network to capture similarities between soils and extend predictions to unseen domains. In practice, this yields slightly reduced accuracy on training soils compared to the one-hot variant, but with superior generalization capacity—a property of particular importance in real-world scenarios where soil information is sparse or incomplete. This balance between accuracy and adaptability highlights embeddings as a promising approach for building scalable, soil-aware PINNs [6].

Nevertheless, embeddings also introduce challenges. Their effectiveness depends on the diversity and representativeness of the training set, and the latent space may absorb not only soil-specific effects but also artifacts from imperfect inputs (e.g., numerical noise in $\theta(z,t)$). This complicates interpretability and calls for careful design, validation, and potentially the integration of additional physical constraints to regularize the learned representations.

In summary, one-hot encoding ensures the highest fidelity when soils are well characterized, while embeddings enable extrapolation to geophysically diverse domains. These findings suggest that the encoding strategy is not a minor implementation detail but a fundamental architectural component in physics-informed modeling of heterogeneous porous media.

IV. Conclusion

This work introduced a hybrid modelling framework for simulating the spatiotemporal transport of ^{137}Cs in variably saturated soils by combining Physics-Informed Neural Networks (PINNs) with physically grounded numerical simulations. Richards equation was solved to generate realistic water content $\theta(z,t)$ and velocity fields, which were coupled to an advection–diffusion–reaction model with kinetic sorption and radioactive decay. These physically consistent fields provided the foundation for training soil-aware PINNs.

The results demonstrate that soil encoding strategies have a decisive impact on predictive accuracy and generalization. The baseline model, which lacks soil awareness, fails to capture heterogeneity and produces biased migration patterns. One-hot encoding yields the most accurate predictions for known soils, faithfully reproducing adsorption and retention behavior. Embedding-based models, while incurring a minor accuracy trade-off on seen soils, preserve latent soil features and generalize to unseen domains. This property is particularly relevant for environmental risk assessment and post-accident remediation, where complete soil characterization is often infeasible.

The methodological contribution of this study lies in framing soil encoding as an architectural choice in PINNs for heterogeneous porous media. Beyond its immediate application to radionuclide transport, the proposed framework establishes a foundation for soil-aware PINNs that could be extended to a range of contaminants and hydrogeological conditions. In practice, the ability to couple sparse data with governing PDEs offers a scalable path toward high-fidelity, data-driven contaminant transport models.

Future directions include incorporating uncertainty quantification to provide confidence intervals for environmental forecasts, exploring transfer learning across contaminants and soil types to reduce data requirements. Another promising avenue is the design of coupled or multi-task PINNs in which one network predicts $\theta(z,t)$ while another predicts $C(z,t)$, both constrained by shared physical laws. Such approaches would enable robust modeling even in data-limited settings, bridging the gap between mechanistic simulations and field-scale prediction.

Overall, this work establishes a methodological and conceptual foundation for soil-aware PINNs. By demonstrating how encoding strategies shape both accuracy and generalization; they provide guidance for future applications of physics-informed machine learning in environmental modeling, risk assessment, and remediation planning.

References

[1] D. Hillel, *Environmental Soil Physics*. San Diego, CA: Academic Press, 1998.

[2] 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. [Online]. Available: https://doi.org/10.1016/j.jcp.2018.10.045

[3] T. Bandai and T. A. Ghezzehei, "Physics-Informed Neural Networks With Monotonicity Constraints for Richardson–Richards Equation," *Water Resources Research*, vol. 57, no. 4, e2020WR027642, 2021. [Online]. Available: https://doi.org/10.1029/2020WR027642

[4] Y. Yang and G. Mei, "A Deep Learning-Based Approach for a Numerical Investigation of Soil–Water Vertical Infiltration with Physics-Informed Neural Networks," *Mathematics*, vol. 10, no. 16, p. 2945, 2022. [Online]. Available: https://doi.org/10.3390/math10162945

[5] T. J. Marshall, J. W. Holmes, and C. W. Rose, *Soil Physics*, 3rd ed. Cambridge University Press, 1996.

[6] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Reviews Physics*, vol. 3, pp. 422–440, 2021. [Online]. Available: https://doi.org/10.1038/s42254-021-00314-5