Learning Soil Water Retention Components through End-to-End Differentiable Hybrid Model

Sarem Norouzi **Per Moldrup** **Ben Moseley** **David Robinson** **Dani Or**

Tobias L. Hohenbrink **Budiman Minasny** **Morteza Sadeghi** **Emmanuel Arthur**

Markus Tuller **Mogens H. Greve** **Lis W. de Jonge**

Abstract

Soil physics is complex, and mechanistic models have traditionally used simplifying assumptions to represent complex processes, but these assumptions can bias predictions. However, the increasing availability of high-quality data offers an opportunity to both improve the predictive power of existing models and gain new fundamental physics insights. Here, we propose a hybrid soil physics framework that combines analytical formulations with flexible, data-driven components to learn uncertain parts directly from data. A key enabler is end-to-end differentiability via automatic differentiation, which allows the entire model, including physical and neural components, to be optimized jointly by minimizing a downstream loss function. We apply this approach to the challenge of partitioning the soil water retention curve (SWRC) into capillary and adsorbed water components. The hybrid model, trained on 483 undisturbed soils from Central Europe, produces smooth and physically consistent SWRC curves and automatically discovers the capillary and adsorptive branches of the curve. Notably, the model reveals a distinctly nonlinear transition between capillary and adsorbed domains, challenging the linear assumptions invoked in previous studies. The methodology introduced here provides a blueprint for learning other soil processes where high-quality datasets are available but mechanistic understanding is incomplete.

1 Introduction

Soils control how water moves, is stored, and becomes available to plants, yet the physics describing these processes is only partially understood. Mechanistic soil models derive equations from physical laws but inevitably simplify complex components, for example, by assuming oversimplified pore geometries or fixed functional forms for soil pore space distribution. While such assumptions make problems tractable, they also bias predictions [1–3].

We propose *differentiable hybrid modeling* as an alternative approach for modeling soil physical processes. Hybrid methods embed neural networks within physical models so that the unknown or poorly understood components of a system can be learned directly from data while the well-established physical laws remain explicitly enforced [4–6]. This combination improves the interpretability and consistency of traditional formulations while enabling discovery of processes that are otherwise inaccessible through purely mechanistic or purely empirical approaches.

We demonstrate this framework on the soil water retention curve (SWRC), which describes how much water a soil holds at different suctions. The SWRC is governed by capillarity and adsorption mechanisms, and separating their contributions is a long-standing challenge in soil physics. Existing models rely on prior assumptions about soil pore geometry and functional forms to describe these

1st Workshop on Differentiable Systems and Scientific Machine Learning @ EurIPS 2025.

components and their transition, but such assumptions strongly affect the partitioning, often producing divergent results even for the same soil [7–10].

Here, we develop a hybrid model that learns the shape of the SWRC and discovers capillary and adsorbed water content components from basic soil properties and data without assuming specific shapes for soil pores or their distribution. The hybrid model couples analytical models for the well-understood parts of the process with constrained neural networks for learning the complex parts. By end-to-end training with automatic differentiation, the model learns the wet end, the capillary–adsorbed transition, and soil-specific parameters. The final SWRC model remains continuous, differentiable, and physically consistent, making it suitable for modeling soil water dynamics.

2 Methods

2.1 Proposed hybrid model

Capillary water refers to liquid water filling the spaces between soil particles, held by surface tension and the contact angle of water with solid surfaces, which leads to the formation of curved liquid–vapor interfaces (menisci). The adsorbed film water component refers specifically to liquid water retained in thin films by adsorptive forces. The total water content can be expressed as the sum of capillary (θ_c) and adsorbed film (θ_a) components:

$$\theta = \theta_c + \theta_a \tag{1}$$

All terms in Eq. (1) are functions of soil suction (*i.e.*, pF). The SWRC at the dry end (*i.e.*, high suctions) becomes linear in pF–θ space, which can be described analytically by the Campbell–Shiozawa model (denoted as θ_{cmp} in Eq. (2)) [11]. In the lower range of pF values, where capillary water begins to contribute, the expression for θ_{cmp} no longer holds. To account for this, we introduce a transition function, denoted as f, that modifies θ_{cmp} in this mixed region. This function is treated as an unknown to be learned from data, and it is expressed as a function of capillary saturation, defined as $S_c = \theta_c/\theta_s$, where θ_s is the saturated water content. Replacing the Campbell–Shiozawa model and the transition function into Eq. (1) yields:

$$\theta = \theta_c + f(S_c)\theta_{\text{cmp}} = \theta_c + f(S_c)\left(1 - \frac{pF}{pF_{\text{dry}}}\right)\theta_o \tag{2}$$

where θ_o and pF_{dry} are fitting parameters of Campbell–Shiozawa model. The parameter pF_{dry} corresponds to the soil suction at oven dryness, where the soil is assumed to reach zero water content.

2.2 Neural networks for learning unknown components

In conventional models, a certain set of rigid assumptions is used to simplify Eq. (2). For instance, the transition function, $f(S_c)$, is assumed as a linearly decreasing function and the capillary component, θ_c, is replaced with a fixed-form sigmoidal parametric function. To avoid these limitations, we replace θ_c with a dedicated neural network, NN_c, which learns the capillary water content as a function of soil basic properties (i.e., sand, silt, clay, organic carbon, bulk density) and pF. Similarly, we replace the transition function, $f(S_c)$, with a neural network that takes in S_c and outputs a scaler that modifies θ_{cmp} for lower suctions. The soil specific constants, θ_s, θ_o, pF_{dry} are likewise replaced with neural networks ($NN_s, NN_o, NN_{\text{dry}}$) that learn these model parameters from soil basic properties.

2.3 Universally accepted physical constraints

We impose a few universally accepted physical constraints on the hybrid ansatz and its interior networks:

1) The capillary component, θ_c, is negligible for $pF > 5$. 2) Total water content remains constant for $pF < -0.3$. 3) To ensure convergence to the Campbell–Shiozawa model at the dry end, the transition function, $f(S_c)$, is hard-constrained to satisfy the condition $f(S_c = 0) = 1$ by reparametrizing this function as: $f(S_c) = 1 + S_c \cdot g(S_c)$, where $g(S_c)$ is a learnable function and is approximated with a neural network, NN_a. 4) The capillary component, θ_c, and θ_o of the Campbell–Shiozawa model must remain below the saturated water content, θ_s. The highest suction at zero water content, pF_{dry}, is constrained to the experimentally observed range of [6.2, 7.6].

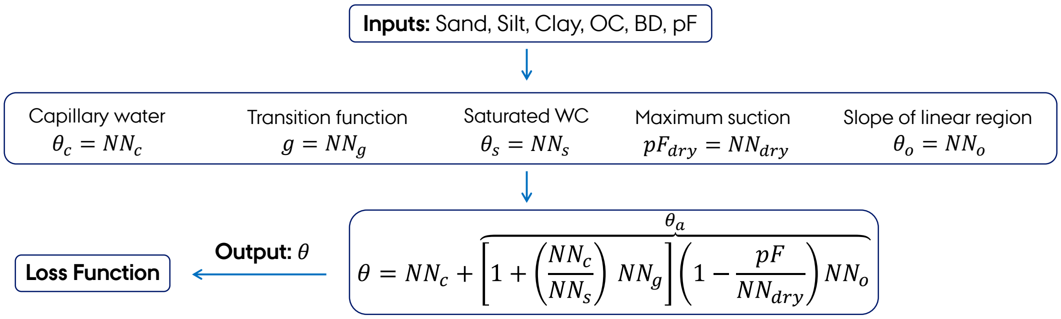

Figure 1: Schematic of the differentiable hybrid model. BD and OC refer to bulk density and organic carbon, respectively. The interior neural networks have no labeled data and are inferred implicitly during end-to-end training.

These constraints are implemented by adding penalizing terms to the loss function, as well as by using specific activation functions in the output layer of different neural networks to limit soil constants to their specified range. All neural networks for estimation of unknown components for soil constants have two hidden layers each with 4 units, except NN_a which was designed slightly more flexible with 16 hidden units in each layer. We used sigmoid activation functions in the output layer of NN_c, NN_s, NN_{dry}, and NN_o. To avoid overfitting, early stopping with a patience of 10 epochs was applied.

2.4 Training with automatic differentiation

The hybrid model developed in this study involves several interconnected neural networks, each containing trainable parameters (Fig. 1). These networks are coupled through a physics-informed ansatz in Eq. (2) that takes in basic soil properties and outputs total water content, θ. For model training, we need the gradients of the loss function with respect to all trainable parameters in each network. We leverage automatic differentiation (AD) [12]. AD automatically constructs a computational graph during the model's forward pass and traces the sequence of mathematical operations from inputs to outputs. During backpropagation, reverse-mode AD traverses this graph from the output layer back to the inputs, systematically applying the chain rule to compute exact gradients with respect to every trainable parameter.

3 Results

We trained the model on SWRC data measured from 483 undisturbed soils in Central Europe, covering a wide range of soil textures and organic carbon contents [13]. The training data consist of water content versus soil suction (pF) without explicit data for the intermediate neural networks. Figure 2 shows the predicted SWRC for six soil samples representing different texture classes. Unlike parametric PTFs, which rely on predefined analytical forms for the SWRC, our hybrid model learns the curve shape directly from the data. Once trained, we predict the entire continuous SWRC by fixing physical properties for each soil sample and varying the pF over a specified range.

The discovered shapes of the SWRCs are smooth, differentiable, and therefore suitable for simulation of soil water flow (i.e., Richardson–Richards equation). These curves exhibit a sigmoidal shape in the wet range and transition to a linear form at lower water contents, consistent with the Campbell–Shiozawa model behavior assumed at the dry end. Notably, the transition between the neural network–predicted region and the analytically modeled region governed by the Campbell–Shiozawa model is seamless, with no noticeable discontinuities or abrupt changes.

The discovered partitioning of capillary and adsorbed film components of the SWRC for the soil are obtained by plotting the first and second terms on the right-hand side of Eq. (2). This data-driven partitioning aligns remarkably well with the physics-based models in the literature that were developed by incorporating detailed interfacial physics within an angular pore geometry [7]. Specifically, the capillary component dominates under saturated conditions for all soils. As pF increases (more suction in soil), pores of varying sizes begin to drain, and this process starts with larger pores. As drainage progresses, water films begin to form along the surfaces of the partially emptied pores. With

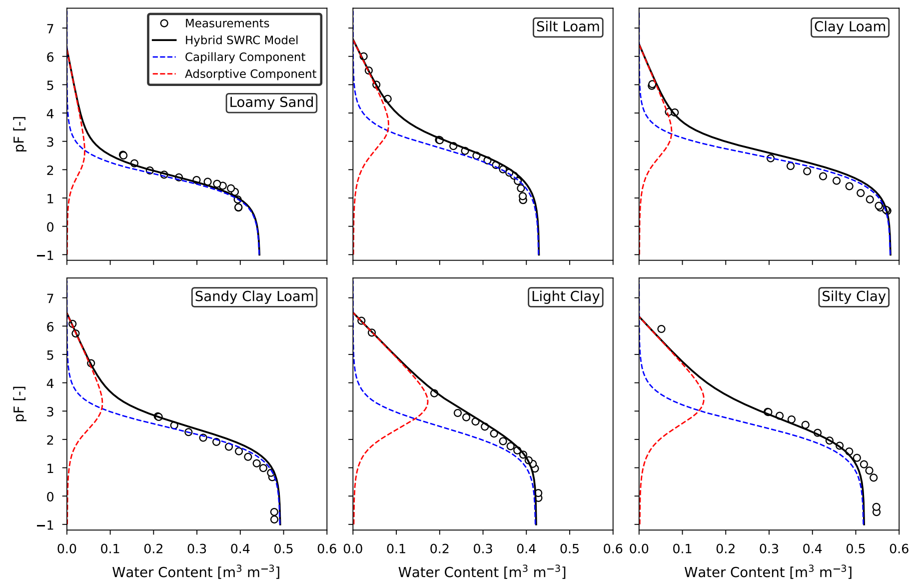

Figure 2: Predicted soil water retention curves for six soil samples with different texture classes.

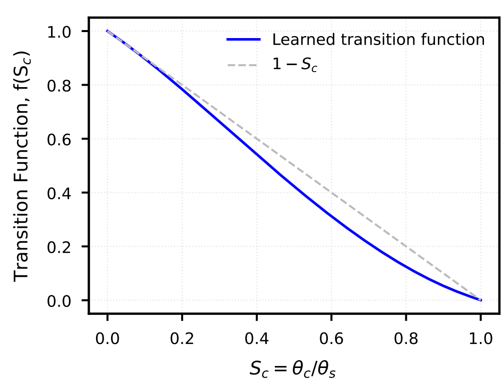

Figure 3: Learned transition function, $f(S_c)$, versus commonly assumed linear model in literature

further increase of pF, smaller pores also undergo drainage, leading to a gradual decrease in the capillary component and a concurrent increase in the contribution of the film component. The water content at the crossover between the capillary and adsorbed components increases for finer textured soils due to the increase in surface area (Fig. 2). Furthermore, the learned transition function, $f(S_c)$, exhibits a distinctly nonlinear behavior, deviating from the commonly assumed linear transition used in previous studies [8, 14] (Fig. 3). Validation of newly discovered soil water retention patterns beyond the training data remains a topic for future research.

4 Conclusion

We proposed a differentiable hybrid model for the soil water retention curve (SWRC) that couples an analytical expression for the dry end with neural networks that learn the unknown components, such as the capillary domain and the transition between regions. The model learns both the overall curve shape and the capillary–adsorbed water partitioning from data, guided by a set of universally accepted physical constraints. Its outputs are in good agreement with predictions from physically based models, yet the approach offers greater flexibility owing to its hybrid nature. Beyond soil water retention, this differentiable framework can be extended to a wide range of soil science problems where process understanding is incomplete but high-quality measurements are available.

References

[1] W. Heber Green and G. A. Ampt. Studies on Soil Phyics. *The Journal of Agricultural Science*, 4(1):1–24, May 1911. ISSN 1469-5146, 0021-8596. doi: 10.1017/S0021859600001441. URL `https://www.cambridge.org/core/journals/journal-of-agricultural-science/article/abs/studies-on-soil-phyics/6EE03D61E70FCEFD6EAE4D59BFCC1FF9`.

[2] WR Gardner. Some steady-state solutions of the unsaturated moisture flow equation with application to evaporation from a water table. *Soil science*, 85(4):228–232, 1958. ISSN 0038-075X. Publisher: LWW.

[3] M. Th. van Genuchten. A Closed-form Equation for Predicting the Hydraulic Conductivity of Unsaturated Soils. *Soil Science Society of America Journal*, 44(5):892–898, 1980. ISSN 1435-0661. doi: 10.2136/sssaj1980.03615995004400050002x. URL `https://onlinelibrary.wiley.com/doi/abs/10.2136/sssaj1980.03615995004400050002x`. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.2136/sssaj1980.03615995004400050002x.

[4] George Karniadakis, Yannis Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, pages 1–19, May 2021. doi: 10.1038/s42254-021-00314-5.

[5] Benjamin Moseley. *Physics-informed machine learning: from concepts to real-world applications*. PhD Thesis, University of Oxford, 2022. URL `https://ora.ox.ac.uk/objects/uuid:b790477c-771f-4926-99c6-d2f9d248cb23`.

[6] Chaopeng Shen, Alison P. Appling, Pierre Gentine, Toshiyuki Bandai, Hoshin Gupta, Alexandre Tartakovsky, Marco Baity-Jesi, Fabrizio Fenicia, Daniel Kifer, Li Li, Xiaofeng Liu, Wei Ren, Yi Zheng, Ciaran J. Harman, Martyn Clark, Matthew Farthing, Dapeng Feng, Praveen Kumar, Doaa Aboelyazeed, Farshid Rahmani, Yalan Song, Hylke E. Beck, Tadd Bindas, Dipankar Dwivedi, Kuai Fang, Marvin Höge, Chris Rackauckas, Binayak Mohanty, Tirthankar Roy, Chonggang Xu, and Kathryn Lawson. Differentiable modelling to unify machine learning and physical models for geosciences. *Nature Reviews Earth & Environment*, 4(8):552–567, August 2023. ISSN 2662-138X. doi: 10.1038/s43017-023-00450-9. URL `https://www.nature.com/articles/s43017-023-00450-9`. Number: 8 Publisher: Nature Publishing Group.

[7] Dani Or and Markus Tuller. Liquid retention and interfacial area in variably saturated porous media: Upscaling from single-pore to sample-scale model. *Water Resources Research*, 35(12): 3591–3605, December 1999. ISSN 0043-1397, 1944-7973. doi: 10.1029/1999WR900262. URL `https://agupubs.onlinelibrary.wiley.com/doi/10.1029/1999WR900262`.

[8] Marc Lebeau and Jean-Marie Konrad. A new capillary and thin film flow model for predicting the hydraulic conductivity of unsaturated porous media. *Water Resources Research*, 46(12), 2010. ISSN 1944-7973. doi: 10.1029/2010WR009092. URL `https://onlinelibrary.wiley.com/doi/abs/10.1029/2010WR009092`. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2010WR009092.

[9] Tobias K. D. Weber, Wolfgang Durner, Thilo Streck, and Efstathios Diamantopoulos. A Modular Framework for Modeling Unsaturated Soil Hydraulic Properties Over the Full Moisture Range. *Water Resources Research*, 55(6):4994–5011, 2019. ISSN 1944-7973. doi: 10.1029/2018WR024584. URL `https://onlinelibrary.wiley.com/doi/abs/10.1029/2018WR024584`. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2018WR024584.

[10] Asghar Ghorbani, Ebrahim Babaeian, Morteza Sadeghi, Wolfgang Durner, Scott B. Jones, and Martinus Th. van Genuchten. An improved van Genuchten soil water characteristic model to account for surface adsorptive forces. *Journal of Hydrology*, page 133692, June 2025. ISSN 0022-1694. doi: 10.1016/j.jhydrol.2025.133692. URL `https://www.sciencedirect.com/science/article/pii/S0022169425010303`.

[11] G. S. Campbell and S. Shiozawa. Prediction of hydraulic properties of soils using particle-size distribution and bulk density data. In *International workshop on indirect methods for estimating the hydraulic properties of unsaturated soils. California: University of California*, pages 317–28, 1992.

[12] Atilim Gunes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: a survey, February 2018. URL http://arxiv.org/abs/1502.05767. arXiv:1502.05767 [cs, stat].

[13] Tobias L. Hohenbrink, Conrad Jackisch, Wolfgang Durner, Kai Germer, Sascha C. Iden, Janis Kreiselmeier, Frederic Leuther, Johanna C. Metzger, Mahyar Naseri, and Andre Peters. Soil water retention and hydraulic conductivity measured in a wide saturation range. *Earth System Science Data*, 15(10):4417–4432, October 2023. ISSN 1866-3516. doi: 10.5194/essd-15-4417-2023. URL https://essd.copernicus.org/articles/15/4417/2023/.

[14] Michael J. Fayer and C. Steven Simmons. Modified Soil Water Retention Functions for All Matric Suctions. *Water Resources Research*, 31(5):1233–1238, 1995. ISSN 1944-7973. doi: 10.1029/95WR00173. URL https://onlinelibrary.wiley.com/doi/abs/10.1029/95WR00173. _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/95WR00173.