22nd International Conference on Sustainable Energy Technologies
Kuala Lumpur, Malaysia, 13 to 15 August 2026 · ISBN 978-1-995850-00-9 · ISSN 2978-4387

CONFERENCE PAPER · AI & RENEWABLE SYSTEMS

An integrated surrogate and optimisation framework for PCM-enhanced energy piles using energy-consistency-enhanced deep operator learning

Chaokai Zhang[1,2], Zhao Xu[1*], Ming Xie[1], Qiliang Yang[2], Zhengheng Gan[1]

[1] School of Civil Engineering, Southeast University, Nanjing 211189, China
[2] School of Defense Engineering, Army Engineering University of PLA, Nanjing 210007, China

RECEIVED 9 February 2026 · REVISED 10 March 2026 · ACCEPTED 18 March 2026
PRESENTED 13 to 15 August 2026 · PUBLISHED 23 November 2026

ABSTRACT

Phase change material (PCM)-enhanced energy piles can improve subsurface heat exchange through latent-heat storage, but their nonlinear thermo-hydraulic response makes direct high-fidelity optimisation computationally expensive. This paper presents an integrated surrogate and optimisation framework based on energy-consistency-enhanced DeepONet (EC-DeepONet). A three-dimensional transient COMSOL model was used to simulate PCM-enhanced energy piles under fixed geometry and operating conditions. From 200 attempted simulations over a 10-dimensional PCM thermophysical parameter space, 134 valid cases were retained after excluding numerically unstable phase-change solutions. Each case was converted into a continuous-response sample containing the outlet-temperature trajectory and cumulative heat exchange over 1–100 h. EC-DeepONet couples interaction-enhanced operator learning with an energy-consistency constraint, enabling joint learning of local temperature response and integrated energy performance. On the validation set, it reduced the outlet-temperature RMSE to 0.0508 °C and the cumulative-energy RMSE to 9.87 MJ, corresponding to reductions of 20.6% and 17.8% relative to standard DeepONet. A five-member EC-DeepONet ensemble was then used to screen 25,000 candidate PCMs generated by Latin hypercube sampling. High-fidelity verification showed that the best candidate achieved 436.51 MJ, exceeding the best original dataset sample by 9.61 MJ. The screening-verification loop required only 3.25 min, compared with 360.64 h for exhaustive high-fidelity scanning. These results demonstrate that energy-consistency-enhanced deep operator learning provides a practical route for rapid PCM pre-selection and uncertainty-aware design screening in energy-pile applications.

KEYWORDS Energy pile, Phase change material, Operator learning, Surrogate model, Optimisation

1. Introduction

Energy piles combine foundation elements with ground heat exchange and are increasingly viewed as an attractive technology for low-carbon buildings. However, the thermal storage capacity of conventional pile materials remains limited, especially under strongly transient operating conditions. To address this limitation, phase change materials (PCMs) have been introduced into energy piles in several forms, including PCM backfill (Cao *et al.* 2022), phase-change concrete (Yang *et al.* 2023), MicroPCM-enhanced concrete (Han *et al.* 2022; Chen *et al.* 2024), and embedded PCM units or tubes (Al Hajaj and Saghir 2023). Experimental and numerical studies have shown that, when the melting range is properly matched to the operating temperature window, PCM integration can improve transient heat exchange (Bao *et al.* 2022; Mousa *et al.* 2022; Gan *et al.* 2025) and can also alter the thermo-mechanical response of energy piles (Liu *et al.* 2025; You *et al.* 2025). Recent work has further indicated that the achieved benefit depends not only on the integration route (Qu *et al.* 2022),

* Corresponding author. Email address: xuzhao@seu.edu.cn (Zhao Xu)

© 2026 The Author(s).
Published by Caravel Press.

doi.org/10.65582/set.2026.028

but also on the coordinated selection of latent heat, phase-change temperature, and PCM content (Zhu *et al.* 2025).

Despite this promise, PCM-enhanced energy piles remain difficult to optimise. High-fidelity models in this study refer to three-dimensional transient numerical models that resolve three-dimensional transient heat transfer, temperature-dependent properties, latent-heat release, and coupled boundary effects. As a result, parameter scanning is both computationally expensive and numerically fragile. This is particularly problematic for engineering design, where optimisation requires repeated evaluation of many thermophysical combinations rather than only a few manually selected cases. Recent data-driven studies on energy piles have demonstrated the value of surrogate modelling for heat transfer per unit length prediction (Dong *et al.* 2026), steady thermal performance estimation (Hu *et al.* 2024), outlet temperature prediction at selected conditions (Chen *et al.* 2023; Wang *et al.* 2025), and long-term displacement assessment (Pei *et al.* 2022). For PCM-enhanced energy piles, however, such scalar-output surrogates remain insufficient because both the transient temperature trajectory and its time-integrated energy consequence are central to engineering evaluation and screening.

The present study investigated whether a continuous-response surrogate could be trained from a limited but physically meaningful COMSOL dataset, whether an energy-consistency constraint could improve cumulative heat-exchange prediction, and whether an ensemble surrogate could support conservative screening of a large PCM design space. To address these questions, an energy-consistency-enhanced DeepONet (EC-DeepONet) framework was developed from the DeepONet operator-learning concept (Lu *et al.* 2021), and then extended into an uncertainty-aware ensemble for material screening.

2. Methods

2.1. High-fidelity numerical model and dataset

The high-fidelity dataset was generated with a three-dimensional transient COMSOL model of a PCM-enhanced energy pile. As summarised in Figure 1, the proposed framework integrates high-fidelity simulation, operator-learning-based surrogate construction, ensemble-assisted screening, and high-fidelity re-evaluation into a closed optimisation loop. Starting from a 10-dimensional PCM thermophysical parameter space, each sampled material combination was embedded into the transient numerical model to generate the corresponding outlet-temperature continuous response, instantaneous heat-transfer rate, and cumulative heat exchange. These outputs were then organised into a PCM temperature characteristic curve (PCM-TCC) dataset, which provides the supervised basis for learning the mapping from material properties to time-dependent thermal behaviour.

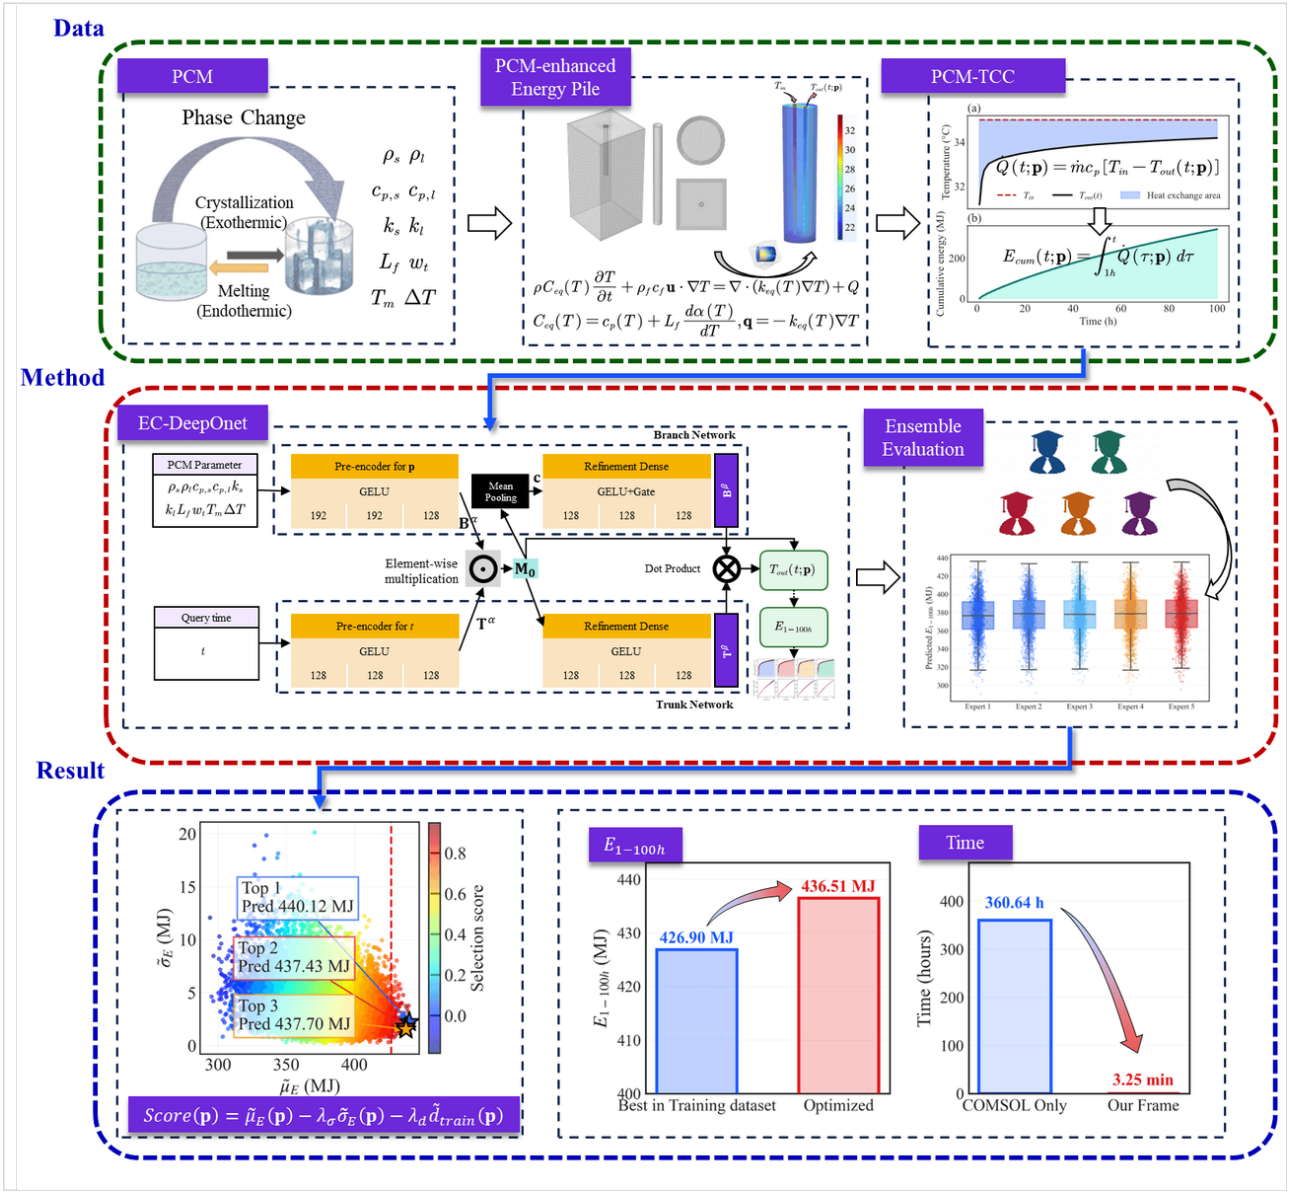

Figure 1. Overall workflow of the proposed framework, including high-fidelity simulation, EC-DeepONet training, ensemble screening, and high-fidelity re-evaluation.

In the surrogate-modelling stage, the PCM parameter vector and query time were used as the joint input of EC-DeepONet to learn the operator mapping from material descriptors to continuous thermal response. The Branch network encoded the thermophysical parameters, while the Trunk network represented the temporal coordinate. Their coupled latent representations were further refined through interaction-enhanced feature extraction so that the outlet-temperature trajectory could be reconstructed over the full time horizon. In addition, an energy-consistency constraint was introduced during training to improve the coordination between local temperature prediction and time-integrated heat-exchange estimation, thereby enhancing the physical consistency of the surrogate for engineering evaluation.

In the screening and optimisation stage, multiple independently trained EC-DeepONet models were assembled into an ensemble evaluator for rapid assessment of large candidate sets. For each candidate PCM, the final screening score combined the ensemble predictive mean, predictive uncertainty, and the distance to the training-sample distribution, so as to favour high-performing yet reliable candidates while reducing extrapolation risk. The top-ranked candidates were then re-evaluated using the high-fidelity COMSOL model, and the best verified design was identified from

this screening-verification loop. In this way, the framework unified material-performance improvement and substantial computational acceleration within a single workflow.

As shown in Figure 2, the fixed physical domain comprised the surrounding soil, pile shaft, and embedded heat-exchange tube, while the geometric configuration and operating conditions were kept unchanged across all simulations. The computational domain measured 10 × 10 × 20 m, while the pile had a length of 10 m and a diameter of 0.8 m. The heat-transfer fluid entered at a constant temperature of 35 °C with a velocity of 0.6 m/s. PCM phase transition was modelled using the apparent heat capacity method, a standard fixed-grid treatment for latent-heat transport in PCM systems (Al-Saadi and Zhai 2015; Khattari *et al.* 2020). In the present study, this allowed the outlet-temperature trajectory and the cumulative heat-exchange response to be extracted consistently from each numerical case.

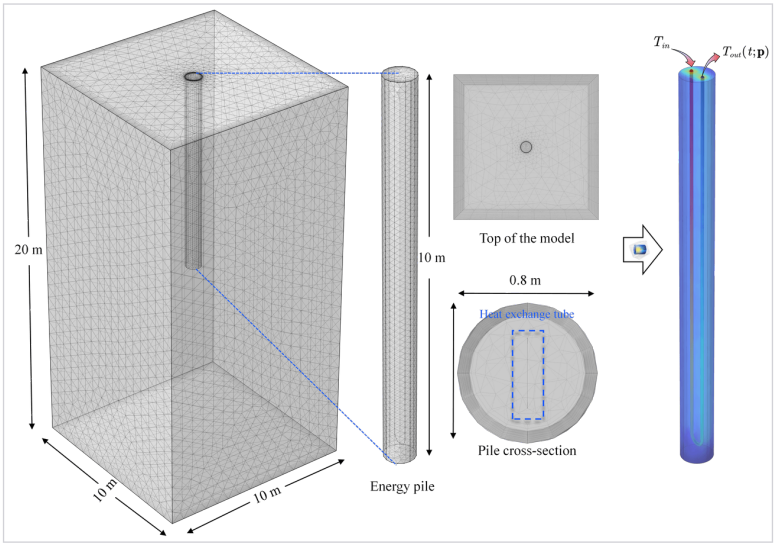

Figure 2. Schematic illustration of the numerical model, including the energy-pile geometry and heat-transfer pathways.

The PCM design vector contained 10 thermophysical descriptors: solid/liquid density, solid/liquid specific heat capacity, solid/liquid thermal conductivity, latent heat of fusion, PCM mass fraction, phase-change temperature, and phase-change interval. To preserve basic thermophysical consistency, inequality constraints such as $\rho_s > \rho_l$, $c_{p,l} > c_{p,s}$, and $k_s \geq k_l$ were enforced during sampling. Although 200 simulations were launched in batch mode, only 134 cases converged successfully because of numerical instability in the nonlinear phase-change regime. These valid cases formed the final supervised dataset.

Each converged case was converted into a continuous-response sample. The outlet-temperature curve was sampled at 100 temporal points over 1–100 h, and the cumulative heat exchange E_{1-100h} was obtained by integrating the instantaneous heat-transfer rate. The dataset was then split at the case level into 70% training and 30% validation subsets so that all temporal points belonging to one physical case remained in the same subset.

Table 1 summarises the sampled PCM parameter space and the thermophysical constraints imposed during high-fidelity case generation.

Table 1. PCM parameter space used for high-fidelity sampling.

Parameter group	Symbols	Range or constraint
Density	ρ_s, ρ_l	$\rho_s \in [800, 950] kg/m^3$, $\rho_l \in [750, 850] kg/m^3$, $\rho_s > \rho_l$
Specific heat capacity	$c_{p,s}, c_{p,l}$	$c_{p,s} \in [1800, 2400]$, $c_{p,l} \in [2000, 2600] J/(kg \cdot K)$, $c_{p,l} > c_{p,s}$
Thermal conductivity	k_s, k_l	$k_s \in [0.20, 0.25]$, $k_l \in [0.14, 0.20] W/(m \cdot K)$, $k_s \geq k_l$
Latent heat	L_f	$[150, 250] kJ/kg$
PCM mass fraction	w_t	$[0.10, 0.25]$
Phase-change temperature	T_m	$[20, 25]^\circ C$
Phase-change interval	ΔT	$[2, 7] K$

2.2. EC-DeepONet surrogate and ensemble screening

Three surrogate models were compared: a multilayer perceptron (MLP), a standard DeepONet, and the proposed EC-DeepONet. The MLP treated each time instant independently by regressing outlet temperature directly from the concatenated input $[\mathbf{p}, t]$. In contrast, DeepONet represented the problem as operator learning, in which a Branch network encoded the PCM parameter vector and a Trunk network encoded time. This formulation was more suitable for continuous-response prediction because it separated operating time from material descriptors and reconstructed the entire curve through a shared latent representation.

EC-DeepONet extended the standard DeepONet in two main ways. First, it introduced an interaction-enhanced latent field before final reconstruction, so that the dependence between material properties and temporal evolution was captured earlier and more explicitly. Second, it imposed an energy-consistency constraint during fine-tuning. The surrogate still predicted outlet temperature as the primary continuous response, but its training objective also penalised the mismatch between the cumulative heat exchange derived from the predicted curve and the corresponding high-fidelity target. This strategy improved the physical consistency between pointwise temperature reconstruction and macro-level energy evaluation.

The resulting architecture, shown in Figure 3, followed the sequence of initial encoding, interaction-field construction, context aggregation, gated refinement, and energy-guided correction. Compared with the conventional DeepONet, the added context-refinement path helped the model capture nonlinear phase-change behaviour more robustly in both the early transient stage and the later plateau stage.

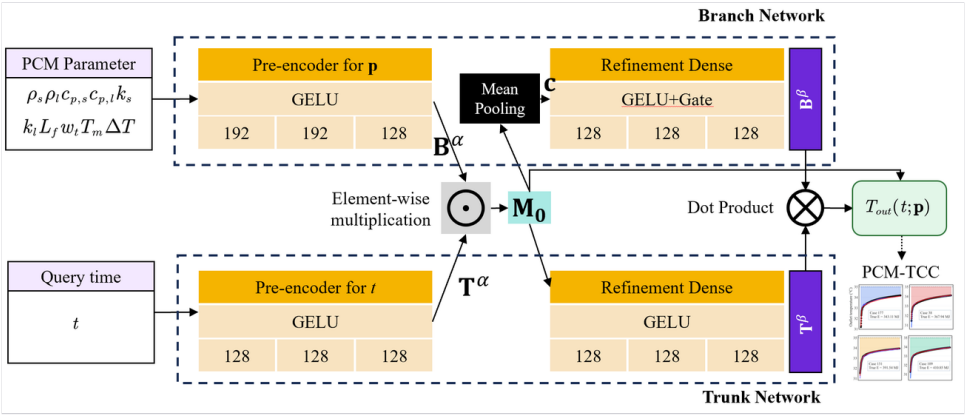

Figure 3. Simplified architecture of EC-DeepONet, highlighting interaction-field construction, context aggregation, gated refinement, and the cumulative-energy consistency pathway.

For optimisation, five EC-DeepONet models with different random initialisations were trained independently and assembled into an ensemble. Candidate PCMs were generated by Latin hypercube

sampling (Stein 1987), and each candidate was scored by combining the ensemble predictive mean, the inter-model standard deviation, and the distance from the training domain. The purpose of this conservative criterion was to reduce the risk of selecting candidates that looked promising only because of extrapolation uncertainty.

3. Results and discussion

3.1. Surrogate accuracy and continuous-response reconstruction

Before surrogate comparison, the structure of the high-fidelity dataset was examined through a parameter-response correlation network based on Spearman rank coefficients (Spearman 1904). As shown in Figure 4, the correlations among the 10 PCM descriptors were generally weak, indicating that the retained valid cases still covered the feasible design space without severe multicollinearity. This point is important for two reasons. First, it suggests that the feasible domain was not collapsed into a narrow manifold after solver failures were removed, so the final 134 valid cases still preserved a meaningful spread of thermophysical combinations. Second, it implies that the subsequent surrogate comparison was not artificially simplified by overly redundant inputs. In other words, the learning task remained genuinely multivariate, which is more representative of practical PCM screening than one-factor sensitivity analysis.

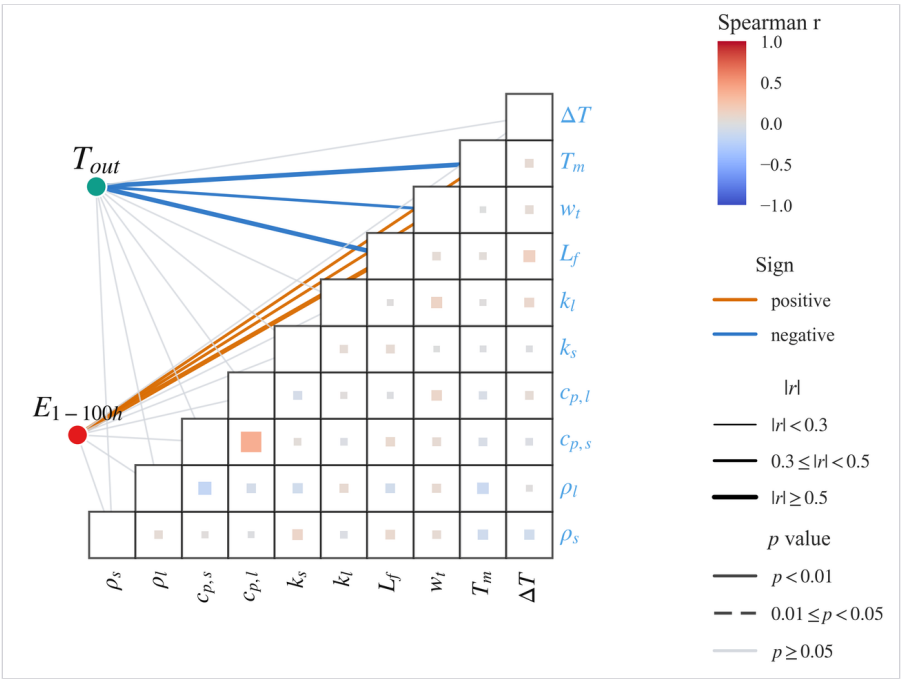

Figure 4. Parameter-response correlation network of the 134 valid high-fidelity cases. Edge colour indicates the sign of the Spearman coefficient, while line weight reflects correlation magnitude and statistical significance.

In contrast, the cumulative heat exchange E_{1-100h} was positively associated with latent heat L_f, PCM mass fraction w_t, and phase-change temperature T_m, whereas the outlet temperature T_{out} showed negative associations with the same variables. This pattern has a clear physical interpretation. Higher latent heat and larger PCM content increase the amount of thermal energy that can be buffered within the pile body, which delays the temperature rise of the circulating fluid and increases the cumulative extracted heat over the analysed period. A suitable upward shift of T_m within the explored range also appears beneficial because it allows the phase-change process to remain active over a larger portion of the operating temperature window considered here. By contrast, parameters such as density and solid-liquid conductivity contrast exhibited weaker isolated

correlations, suggesting that their influence is more conditional and depends on interaction with other descriptors rather than on a simple monotonic trend alone.

Another useful observation from Figure 4 is that no single descriptor dominated the entire response space by itself. Even for variables that showed the strongest associations, the correlation strengths remained moderate rather than near-deterministic. This indicates that the energy-pile response is governed by coordinated parameter matching instead of a single “best” property. From a modelling perspective, such a structure favours continuous-response operator learning: the surrogate must capture coupled effects across material properties and time, rather than interpolate a nearly one-dimensional response surface. The correlation network therefore provides not only a preliminary physical interpretation of the dataset, but also a justification for the modelling strategy adopted in this study.

The three surrogate models were first compared on the validation set. As shown in Table 2, EC-DeepONet provided the best performance in all three aspects considered in this conference paper: outlet-temperature reconstruction, cumulative-energy prediction, and rank preservation across materials. Relative to the standard DeepONet, the temperature RMSE decreased from 0.0639 °C to 0.0508 °C, the energy RMSE decreased from 12.01 MJ to 9.87 MJ, and the Spearman rank correlation increased from 0.911 to 0.925. These improvements indicated that the added interaction and energy-consistency mechanisms were beneficial not only for local curve fitting but also for the decision-relevant ranking task.

Table 2. Validation accuracy of the surrogate models.

Model	Temperature RMSE (°C)	Energy RMSE (MJ)	Spearman rank correlation
MLP	0.0793	12.37	0.908
DeepONet	0.0639	12.01	0.911
EC-DeepONet	0.0508	9.87	0.925

The metric pattern also highlights the difference between pointwise regression and continuous-response operator learning. The MLP produced the largest errors because it treated each time instant as an independent sample and therefore did not explicitly preserve temporal coherence across the full response curve. Standard DeepONet already improved the results by separating the material descriptor space from the temporal coordinate and by reconstructing the response through a shared latent representation. EC-DeepONet extended this advantage further. Its interaction-enhanced latent field enabled stronger coupling between time evolution and material properties, while the additional energy-consistency constraint discouraged temperature trajectories that might look locally acceptable but lead to biased integrated heat exchange. This is particularly relevant for engineering evaluation, because design decisions are rarely based on isolated time points alone.

From an engineering point of view, the absolute error magnitudes are also meaningful. An outlet-temperature RMSE of 0.0508 °C is small relative to the overall response span in the analysed transient process, indicating that the proposed model can reconstruct the thermal curve with limited local distortion. However, the energy RMSE is arguably the more important metric for screening, because optimisation in this study ultimately targets cumulative heat-exchange performance. The fact that the relative improvement in energy prediction is comparable to, and slightly larger than, that in temperature reconstruction suggests that the proposed training strategy did more than simply smooth the curve visually. It improved the physical usefulness of the prediction for material ranking and optimisation.

To examine the continuous-response behaviour more directly, Figure 5 compares four representative validation cases. The MLP showed visible deviations in the early fast-rise stage, whereas the standard DeepONet reproduced the general trend but retained a noticeable bias in the long-term plateau for some cases. EC-DeepONet remained closer to the discrete COMSOL points

throughout both stages. This behaviour was especially important because even small temperature errors may accumulate after temporal integration and lead to larger errors in cumulative energy. The improved agreement between the EC-DeepONet temperature curves and the COMSOL data therefore translated into better energy prediction.

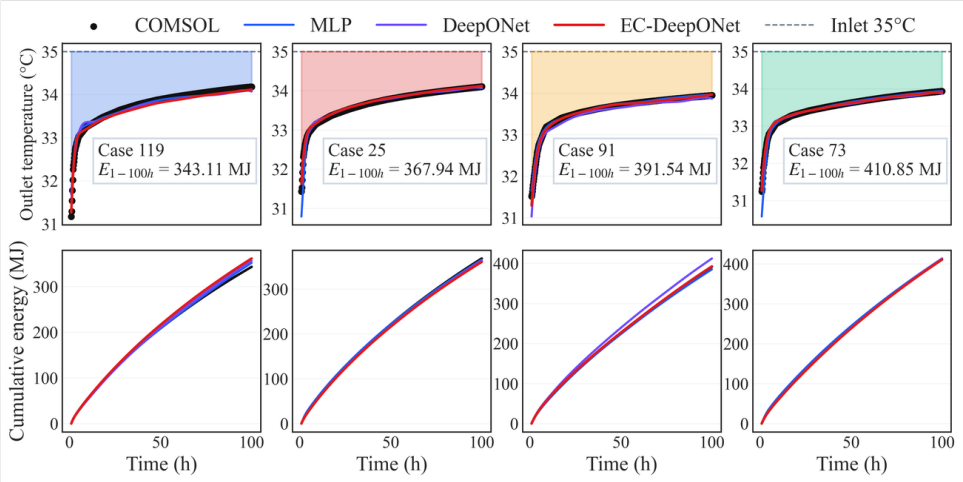

Figure 5. Comparison of outlet-temperature reconstruction and cumulative heat exchange for four representative validation cases. EC-DeepONet provided the closest match to the high-fidelity responses across both the transient and plateau stages.

The case-wise comparison also helps clarify where the difficulty of the prediction task comes from. In the early transient stage, the outlet temperature is controlled mainly by the initial fluid-solid temperature difference, local heat capacity, and the onset of latent-heat absorption. The response can change rapidly over a short period, which explains why the MLP, lacking a global response structure, was more prone to local misalignment. In the later stage, the challenge shifts. Once part of the PCM reservoir has been activated, the response depends increasingly on the duration of the phase-change plateau, the effective thermal resistance of the composite medium, and the degree to which latent heat is gradually exhausted. Standard DeepONet captured the overall curve shape, but in some cases the plateau level or its temporal drift remained slightly biased, which is precisely the type of systematic deviation that can distort integrated heat estimates.

EC-DeepONet showed a more stable behaviour in both regimes. This suggests that the proposed architecture was not merely fitting average trends, but was better at preserving response-shape consistency across materials with different transition characteristics. In practical terms, this means that the model can more reliably distinguish between candidate PCMs whose curves differ by subtle shifts in plateau duration or slope rather than by large differences in peak temperature. Such distinctions are difficult to recover using scalar surrogates or isolated time-point regressors, but they are central to performance assessment in PCM-enhanced thermal systems.

Figure 6 summarises the comparative validation performance of the candidate surrogate models across the three adopted metrics.

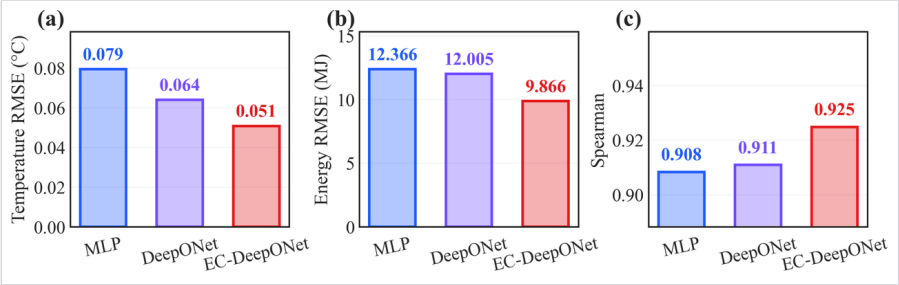

Figure 6. Summary comparison of surrogate-model accuracy. EC-DeepONet achieved the best overall balance between temperature accuracy, energy accuracy, and ranking consistency.

Figure 6 also shows that the performance advantage of EC-DeepONet is balanced rather than metric-specific. In many surrogate studies, an apparent gain in one error measure can be offset by deterioration in ranking stability or by loss of robustness in integrated outputs. That pattern was not observed here. Although the increase in Spearman correlation from 0.911 to 0.925 may appear numerically modest, it is meaningful for top-candidate discovery because optimisation often operates in a regime where many materials have similar predicted benefits. A small improvement in ordering consistency can therefore translate into a substantial improvement in the probability that a high-ranked candidate remains competitive after high-fidelity verification.

3.2. Ensemble screening and optimisation outcome

After validating the single-model surrogate, the five-member EC-DeepONet ensemble was used to screen 25,000 candidate PCM combinations. Among them, 112 candidates were predicted to outperform the best observed sample in the original high-fidelity dataset, whose cumulative heat exchange was 426.90 MJ. High-fidelity verification was then performed on the top-ranked candidates returned by the ensemble screening process.

The proportion of apparently superior candidates was therefore below 0.5% of the screened design pool. This observation is useful because it shows that improved solutions do not occupy the design space broadly; rather, they are concentrated in a comparatively narrow region. Under such circumstances, a surrogate used only for coarse trend identification would not be sufficient. The screening framework must be able to distinguish small but meaningful performance differences while remaining cautious about extrapolation. This is one reason why the ensemble strategy is valuable: it provides not only a mean estimate of performance, but also an uncertainty signal that helps decide which promising candidates are trustworthy enough to justify high-fidelity re-evaluation.

The verification results are summarised in Table 3. A useful observation is that the candidate with the highest predicted mean was not the best after COMSOL re-evaluation. Instead, the second-ranked candidate achieved the highest verified cumulative heat exchange, reaching 436.51 MJ. This candidate exceeded the best observed sample by 9.61 MJ, equivalent to an improvement of about 2.25%. This result supported the rationale of conservative screening: using predictive uncertainty together with predictive benefit can improve the chance of selecting candidates that remain strong after high-fidelity re-evaluation.

Table 3. High-fidelity verification results for the top-ranked candidates selected by the ensemble surrogate.

Candidate rank	Predicted mean (MJ)	Predicted std (MJ)	Verified energy (MJ)	Verification time (s)
1	440.12	2.23	423.84	54.03
2	437.43	1.77	436.51	61.02
3	437.70	1.52	428.55	66.04

This outcome deserves careful interpretation. If selection had been based on predictive mean alone, the first-ranked candidate would have appeared to be the obvious optimum. However, its verification result fell well below its predicted value, whereas the second-ranked candidate proved markedly more reliable. The gap between these two cases reflects the practical reality of surrogate-assisted optimisation in nonlinear thermal systems: the most optimistic prediction is not always the best engineering choice when uncertainty and distribution shift are taken seriously. Candidate 3 provides a complementary perspective. It had a predictive standard deviation smaller than Candidate 1, yet its verified energy was still lower than that of Candidate 2. This means that uncertainty alone is not sufficient as a selection rule; reliable screening requires the joint consideration of benefit, uncertainty, and sample-distribution proximity.

The verification timings in Table 3 also confirm that the role of the high-fidelity model was preserved rather than bypassed. Each re-evaluated case still required roughly one minute of COMSOL computation, which is manageable for a handful of selected candidates but impractical for exhaustive scanning. The surrogate therefore did not replace physics-based simulation as the final authority. Instead, it acted as a front-end accelerator that substantially reduced the number of candidates requiring expensive verification. This division of labour is important from an engineering workflow perspective, because it preserves trust in the final design while dramatically reducing total search cost.

From the perspective of parameter trends, the verified optimum combined high latent heat, a relatively high PCM mass fraction, increased solid/liquid heat capacities, and improved thermal conductivity with a moderately reduced phase-change temperature and a narrowed phase-change interval. The improvement therefore appeared to arise from coordinated parameter matching rather than from pushing a single variable to its upper bound. This interpretation is consistent with earlier numerical work showing that PCM performance depends on the selected enhancement route (Aljabr *et al.* 2021) and with experimental and numerical studies indicating that material properties must be matched to the operating thermal window (Yang *et al.* 2021; Yang *et al.* 2023).

Physically, this combination is reasonable. Higher latent heat and a larger PCM mass fraction directly enlarge the available enthalpy buffer, while higher heat capacities help smooth temperature evolution before and after the main transition interval. Improved conductivity reduces the internal thermal bottleneck that can otherwise prevent the latent-heat reservoir from being effectively mobilised. At the same time, the best design did not simply maximise T_m or broaden ΔT. Instead, a moderately reduced phase-change temperature together with a narrowed transition interval appears to have aligned the phase transition more effectively with the operating temperature trajectory of the pile system, so that latent heat could be activated earlier and more intensively within the analysed time window. This again supports the view that optimisation should target thermophysical coordination rather than isolated property maximisation.

The engineering relevance of the surrogate became clearest when computational effort was considered. Exhaustive high-fidelity evaluation of all 25,000 candidates would have required approximately 360.64 h, while the complete surrogate-screening and verification loop required only 3.25 min. Even though the absolute gain in cumulative heat exchange was moderate, the reduction in design-cycle duration made repeated material screening and iterative design updates far more practical.

Expressed differently, the proposed framework reduced the effective optimisation time by roughly 6,600 times. This scale of acceleration is significant even if the final thermodynamic improvement appears moderate in percentage terms. In practical design, a few percent increase in cumulative heat exchange can still be valuable when it is obtained without repeated manual trial-and-error and without prohibitive computational cost. More importantly, the reduced turnaround time changes what is feasible in the design process itself: instead of evaluating only a few manually selected candidate materials, engineers can examine broad parameter regions, compare multiple

optimisation scenarios, and update assumptions quickly when operating conditions or design priorities change.

At the same time, the present results should be interpreted within the scope of the current study. The surrogate was trained on 134 valid high-fidelity cases under fixed geometry and operating conditions, so the reported optimum is conditional on that design context. For different pile dimensions, inlet conditions, or surrounding-soil properties, the numerical model and training data would need to be updated accordingly. Nevertheless, the results demonstrate the methodological value of the framework: once a representative high-fidelity dataset is available, continuous-response operator learning with uncertainty-aware ensemble screening provides an efficient route for PCM pre-selection and for identifying promising regions of the thermophysical design space.

Figure 7 highlights the combined benefit of the proposed screening framework in terms of verified heat-exchange improvement and optimisation-time reduction.

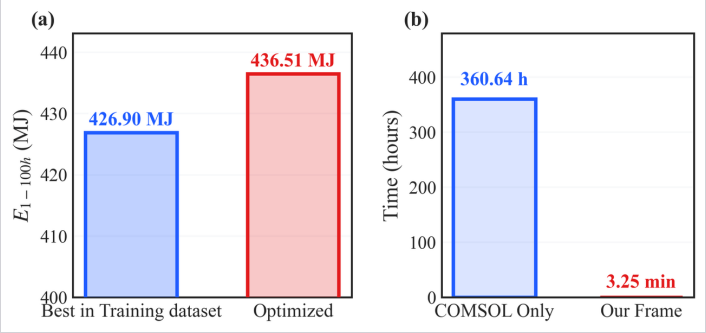

Figure 7. Improvement in the best cumulative heat exchange and reduction in optimisation time achieved by the proposed EC-DeepONet-based screening framework.

4. Conclusions

This paper presented a surrogate-based PCM optimisation framework for energy piles. The main conclusions were threefold.

(1) Continuous-response operator learning was more suitable than pointwise regression for this nonlinear transient problem, and EC-DeepONet achieved the best validation accuracy among the tested models.

(2) The energy-consistency constraint improved not only cumulative-energy prediction but also the reliability of material ranking, which is essential for engineering screening.

(3) The ensemble-guided optimisation workflow identified a verified PCM design with 436.51 MJ cumulative heat exchange while reducing the effective optimisation time from hundreds of hours to a few minutes.

Overall, the results support EC-DeepONet as a practical tool for rapid PCM pre-selection and uncertainty-aware material screening in shallow geothermal systems.

Declarations

Author contributions. Chaokai Zhang: Conceptualization, Methodology, Investigation, Formal analysis, Data curation, Writing (original draft), Writing (review and editing), Visualization, Validation, Software, Funding acquisition. Zhao Xu: Supervision, Project administration, Resources. Ming Xie: Supervision. Qiliang Yang: Supervision, Project administration, Resources. Zhengheng Gan: Conceptualization, Methodology, Investigation, Formal analysis, Data curation.

Funding. This research was funded by the National Natural Science Foundation of China (72671069 and 52178307) and the Postgraduate Research & Practice Innovation Program of Jiangsu Province (26CXJH0600).

Conflicts of interest. The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability. Data will be made available on request.

Generative AI. During the preparation of this work, the authors used Claude to improve language, readability, and editorial presentation. The authors reviewed and edited all generated content and take full responsibility for the final manuscript.

Acknowledgements. This work was supported by the National Key R&D Program of China (Grant No. 2021YFF0500900).

References

Al Hajaj, Zainab; Saghir, M.Z. (2023) 'Numerical study on the influence of embedded PCM tubes on the energy storage properties of the geothermal energy pile', *International Journal of Thermofluids*, 20, pp. 100416. Available at: https://doi.org/10.1016/j.ijft.2023.100416

Al-Saadi, Saleh Nasser; Zhai, Zhiqiang (John) (2015) 'Systematic evaluation of mathematical methods and numerical schemes for modeling PCM-enhanced building enclosure', *Energy and Buildings*, 92, pp. 374-388. Available at: https://doi.org/10.1016/j.enbuild.2015.01.044

Aljabr, Ahmad; Chiasson, Andrew; Alhajjaji, Amr (2021) 'Numerical Modeling of The Effects of Micro-Encapsulated Phase Change Materials Intermixed with Grout in Vertical Borehole Heat Exchangers', *Geothermics*, 96, pp. 102197. Available at: https://doi.org/10.1016/j.geothermics.2021.102197

Bao, Xiaohua; Qi, Xuedong; Cui, Hongzhi; Tang, Waiching; Chen, Xiangsheng (2022) 'Experimental study on thermal response of a PCM energy pile in unsaturated clay', *Renewable Energy*, 185, pp. 790-803. Available at: https://doi.org/10.1016/j.renene.2021.12.062

Cao, Ziming; Zhang, Guozhu; Liu, Yiping; Zhao, Xu (2022) 'Thermal performance analysis and assessment of PCM backfilled precast high-strength concrete energy pile under heating and cooling modes of building', *Applied Thermal Engineering*, 216, pp. 119144. Available at: https://doi.org/10.1016/j.applthermaleng.2022.119144

Chen, Yu; Kong, Gangqiang; Xu, Xiaoliang; Hu, Shuaijun; Yang, Qing (2023) 'Machine-learning-based performance prediction of the energy pile heat pump system', *Journal of Building Engineering*, 77, pp. 107442. Available at: https://doi.org/10.1016/j.jobe.2023.107442

Chen, Kehao; Shen, Yang; Liu, Jiaxin; Wang, Chaoran; Han, Chanjuan (2024) 'Thermomechanical performance analysis of MicroPCM-enhanced energy pile concrete: A multiscale numerical study and formulation design implications', *Journal of Building Engineering*, 86, pp. 108791. Available at: https://doi.org/10.1016/j.jobe.2024.108791

Dong, Siming; Wang, Chenglong; Liu, Hanlong; Bouazza, Abdelmalek; Ding, Xuanming; Kong, Gangqiang (2026) 'A machine learning hybrid model for heat transfer per meter prediction of energy pipe pile', *Geothermics*, 134, pp. 103521. Available at: https://doi.org/10.1016/j.geothermics.2025.103521

Gan, Zhengheng; Yin, Mei; Dou, Shiqi; Chang, Hong; Jiang, Huicheng (2025) 'Enhanced thermal performance of energy piles based on composite phase change material and random aggregate modelling', *Energy*, 341, pp. 139511. Available at: https://doi.org/10.1016/j.energy.2025.139511

Han, Chanjuan; Shen, Yang; Chen, Kehao; Gong, Cheng; Liu, Tianle; Qu, Bo (2022) 'Characteristics and energy performance of novel MicroPCM C50 energy pile in cooling mode', *Energy and Buildings*, 274, pp. 112442. Available at: https://doi.org/10.1016/j.enbuild.2022.112442

Hu, Shuaijun; Kong, Gangqiang; Zhang, Changsen; Fu, Jinghui; Li, Shiyao; Yang, Qing (2024) 'Data-driven models for the steady thermal performance prediction of energy piles optimized by metaheuristic algorithms', *Energy*, 313, pp. 134000. Available at: https://doi.org/10.1016/j.energy.2024.134000

Khattari, Y.; El Rhafiki, T.; Choab, N.; Kousksou, T.; Alaphilippe, M.; Zeraouli, Y. (2020) 'Apparent heat capacity method to investigate heat transfer in a composite phase change material', *Journal of Energy Storage*, 28, pp. 101239. Available at: https://doi.org/10.1016/j.est.2020.101239

Liu, Jiaxin; Wang, Chaoran; Chen, Kehao; Shen, Yang; Han, Chanjuan (2025) 'Comparative experimental study on thermo-mechanical responses of phase-change and normal energy piles', *Geomechanics for Energy and the Environment*, 44, pp. 100758. Available at: https://doi.org/10.1016/j.gete.2025.100758

Lu, Lu; Jin, Pengzhan; Pang, Guofei; Zhang, Zhongqiang; Karniadakis, George Em (2021) 'Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators', *Nature Machine Intelligence*, 3(3), pp. 218-229. Available at: https://doi.org/10.1038/s42256-021-00302-5

Mousa, M.M.; Bayomy, A.M.; Saghir, M.Z. (2022) 'Long-term performance investigation of a GSHP with actual size energy pile with PCM', *Applied Thermal Engineering*, 210, pp. 118381. Available at: https://doi.org/10.1016/j.applthermaleng.2022.118381

Pei, Huafu; Song, Huaibo; Meng, Fanhua; Liu, Weiling (2022) 'Long-term thermomechanical displacement prediction of energy piles using machine learning techniques', *Renewable Energy*, 195, pp. 620-636. Available at: https://doi.org/10.1016/j.renene.2022.06.057

Qu, Bo; Liu, Tianle; Gong, Cheng; Liu, Jiaxin; Shen, Yang; Chen, Kehao; Han, Chanjuan (2022) 'Investigation of microencapsulate phase change material-based energy pile group: Energy analysis and optimization design', *Journal of Cleaner Production*, 381, pp. 135204. Available at: https://doi.org/10.1016/j.jclepro.2022.135204

Spearman, C. (1904) 'The Proof and Measurement of Association between Two Things', *The American Journal of Psychology*, 15(1), pp. 72. Available at: https://doi.org/10.2307/1412159

Stein, Michael (1987) 'Large Sample Properties of Simulations Using Latin Hypercube Sampling', *Technometrics*, 29(2), pp. 143-151. Available at: https://doi.org/10.1080/00401706.1987.10488205

Wang, Chenglong; Dong, Siming; Bouazza, Abdelmalek; Ding, Xuanming (2025) 'Explainable machine learning models to predict outlet water temperature of pipe-type energy pile', *Renewable Energy*, 246, pp. 122972. Available at: https://doi.org/10.1016/j.renene.2025.122972

Yang, Weibo; Sun, Taofu; Yang, Binbin; Wang, Feng (2021) 'Laboratory study on the thermo-mechanical behaviour of a phase change concrete energy pile in summer mode', *Journal of Energy Storage*, 41, pp. 102875. Available at: https://doi.org/10.1016/j.est.2021.102875

Yang, Weibo; Yang, Binbin; Wang, Feng; Jing, Naidong (2023) 'Numerical evaluations on the effects of thermal properties on the thermo-mechanical behaviour of a phase change concrete energy pile', *Energy and Built Environment*, 4(1), pp. 1-12. Available at: https://doi.org/10.1016/j.enbenv.2021.05.008

You, Tian; Yan, Tongrui; Huang, Yiyu; Huang, Shuai; Cui, Hongzhi (2025) 'Thermo-mechanical coupling performance analysis of phase change material and nanofluid synergized high-efficiency energy piles for shallow geothermal energy utilization', *Energy*, 334, pp. 137691. Available at: https://doi.org/10.1016/j.energy.2025.137691

Zhu, Xiang; Chen, Tao; Li, Yunhua; Gao, Xinjun (2025) 'Numerical investigation on the energy efficiency analysis of PCM enhanced PHC energy pile under building cooling', *Journal of Building Engineering*, 107, pp. 112709. Available at: https://doi.org/10.1016/j.jobe.2025.112709

HOW TO CITE THIS PAPER

Zhang, C., Xu, Z., Xie, M., Yang, Q. and Gan, Z. (2026). An integrated surrogate and optimisation framework for PCM-enhanced energy piles using energy-consistency-enhanced deep operator learning. In: *Proceedings of the 22nd International Conference on Sustainable Energy Technologies (SET2026)*, Kuala Lumpur, Malaysia, 13 to 15 August 2026. Caravel Press, pp. 316–328. ISBN 978-1-995850-00-9. ISSN 2978-4387. https://doi.org/10.65582/set.2026.028

COPYRIGHT AND LICENCE

© 2026 The Author(s). This is an open access paper distributed under the terms of the Creative Commons Attribution (CC BY 4.0) licence, which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

WSSET World Society of Sustainable Energy Technologies Organised by the World Society of Sustainable Energy Technologies. Published by Caravel Press.