Automation in Construction 190 (2026) 107099

Contents lists available at ScienceDirect

Automation in Construction

journal homepage: www.elsevier.com/locate/autcon

Physics-constrained and LLM-driven framework for cross-project deformation prediction in deep foundation pits

Yue Pan [a], Zonghang He [a,*], Jin-Jian Chen [a], Xiang Li [b], Jianfeng Zhou [c]

[a] State Key Laboratory of Ocean Engineering, Shanghai Key Laboratory for Digital Maintenance of Buildings and Infrastructure, Department of Civil Engineering, Shanghai Jiao Tong University, 800 Dongchuan Road, Shanghai 200240, China

[b] Research Institute of Advanced Materials (Shenzhen) Co., Ltd, Shenzhen 518045, China

[c] College of Civil Engineering, Huaqiao University, Jimei Avenue 668, Xiamen 361021, China

ARTICLE INFO

Keywords:
Large language model
Domain-adversarial neural network
Spatiotemporal deformation prediction
Deep excavation

ABSTRACT

Data-driven models for deep excavation deformation prediction often lack cross-project transferability and physical consistency. This paper introduces PL-DAST, a physics-constrained, large-language-model (LLM)-driven domain-adversarial spatiotemporal framework. PL-DAST embeds textual geological knowledge into mechanism-regularized learning: an LLM converts geological investigation reports into structured semantic embeddings, adversarial training aligns these representations across projects to learn domain-invariant features, and soil loss theory is incorporated as a mechanistic regularizer to promote physically plausible dynamics. Validated on three metro excavation projects, PL-DAST achieves high predictive accuracy (R^2_{acc}= 0.993) and strong physical consistency (R^2_{phy} = 0.968), while reducing physics loss by 91.52% relative to data-only baselines. The LLM-derived embeddings capture shared stratigraphic signatures (minimum cosine similarity = 0.895) while preserving project-specific variation. Feature alignment further improves transferability, reducing the average sliced Wasserstein distance by 11.81% (from 0.127 to 0.112). Overall, PL-DAST provides a transferable and mechanism-consistent deformation prediction tool for intelligent monitoring and risk-informed decision-making in deep excavation.

1. Introduction

Rapid urbanization is intensifying pressure on land resources, while surface transport networks in many metropolises are approaching capacity [1,2]. Accordingly, underground space is increasingly regarded as a strategic asset for sustainable urban development [3,4]. Deep excavation underpins this expansion, and its construction quality and safety are therefore pivotal to the performance of underground infrastructure [5,6]. However, excavation projects remain vulnerable to major hazards, including support instability and excessive ground settlement [7,8]. Effective risk mitigation requires predictive tools that are not only accurate but also physically credible and transferable across projects with heterogeneous geological conditions. This highlights the need for a deformation prediction framework that integrates data-driven modeling with geotechnical mechanisms and geological text-derived features to deliver predictions that are both accurate and physically consistent.

Deformation prediction is central to safety management in deep excavation engineering [9,10]. Conventional approaches, including the elastic foundation beam method [11,12] and finite element method (FEM) analysis [13,14], are widely used for early-stage predictions but exhibit clear limitations. The elastic foundation beam method relies on oversimplified assumptions that can introduce large errors, whereas FEM offers higher fidelity but depends strongly on parameter calibration and incurs substantial computational cost, limiting routine use [15,16]. Deep learning (DL) has therefore gained increasing attention, owing to its ability to capture nonlinear relationships directly from monitoring data [17,18]. Recent studies have illustrated its potential: Wang et al. [19] applied an Attention-based Ensemble BiLSTM network with ARIMA (Att-EBiLSTM-ARIMA) to predict excavation risk states, achieving R^2 exceeding 0.98. Li et al. [2] used a bidirectional long short-term memory (Bi-LSTM) model to predict excavation-induced risks to adjacent buildings, with MSEs below 0.06. Zhou et al. [20] developed an attention-causality graphical gated network (AC-GGN) for ground settlement prediction, yielding an RMSE of 0.59 in practice. Pan et al. [8] proposed an online learning-based multi-attribute spatiotemporal transformer network (OMSTTN), yielding an R^2 of 0.94 for excavation-

* Corresponding author.
E-mail address: hzh20202765@sjtu.edu.cn (Z. He).

https://doi.org/10.1016/j.autcon.2026.107099
Received 13 January 2026; Received in revised form 15 May 2026; Accepted 5 June 2026
Available online 10 June 2026

induced risk prediction. However, these impressive quantitative metrics are typically highly site-specific, the predictive accuracy of these purely data-driven models often degrades significantly when transferred to new excavation sites due to the absence of domain adaptation and physical constraints.

Despite advances in deep-learning-based deformation prediction, a key limitation remains: purely data-driven models often ignore geotechnical mechanisms, which can lead to violations of governing principles and undermine reliability, particularly when transferred across projects [21,22]. Physics-informed neural networks (PINNs) address this weakness by embedding physical laws into training through physics-based loss terms [23,24]. This preserves the strong feature-extraction capabilities of deep learning while constraining predictions to remain physically consistent, thereby enhancing interpretability and generalization under data scarcity or distribution shifts [25,26]. Recent studies highlight the advantages of PINNs. For example, Li et al. [27] incorporated a fuzzy partial differential equation derived from fatigue cracking theory into a PINN for pavement performance prediction, improving physical consistency. Yan et al. [28] developed a PINN-based bridge-bearing control strategy that limited force adjustment errors to within 18% and outperformed 20 benchmark machine-learning models. Furthermore, recent studies demonstrate the versatility of PINNs by enforcing distinct physical laws through two primary strategies: explicitly embedding physical invariance conditions directly into the network architecture [29], and integrating complex governing equations alongside advanced constitutive models, including linear elasticity [30], von Mises plasticity [31], and non-associative Drucker-Prager flow rules [32], into the loss function. Ultimately, both approaches significantly improve solution accuracy and robustness across diverse mechanics applications. These results conclusively prove that embedding physical mechanisms can yield more physically credible predictions and improve robustness across diverse engineering contexts. This therefore supports the development of physics-constrained and transferable deformation prediction frameworks for deep foundation pits.

In addition, textual data, such as geological investigation reports and construction logs, contain rich semantic information that is critical for interpreting deformation in deep excavations. However, most existing prediction models make limited use of these unstructured records, because conventional DL architectures are not designed to parse and integrate free text effectively. As a result, valuable geological context and engineering judgement are frequently lost. Large language models (LLMs) offer a promising solution. Pre-trained on large scale corpora, LLMs can capture contextual semantics, infer implicit geological attributes and encode engineering knowledge expressed in natural language [33,34]. By transforming unstructured text into structured semantic embeddings, they enable direct fusion of textual and numerical inputs. Crucially, LLMs' capacity to generalize across diverse writing styles and previously unseen expressions supports more comprehensive use of textual evidence, enriching feature representations and improving the accuracy, interpretability and reliability of deformation prediction.

Although PINNs can enforce physical consistency and LLMs can exploit textual evidence, deformation prediction in deep excavations remains limited by poor generalization across heterogeneous engineering settings [35,36]. This limitation arises largely from inter-project differences in geology and construction conditions, which often cause substantial performance degradation when models are deployed at new sites [37]. To mitigate cross-project distribution shifts, this study introduces Domain-Adversarial Neural Networks (DANNs) [38]. DANNs train feature extractors and domain discriminators in opposition, suppressing project-specific signals and yielding domain-invariant features (DIFs) [39,40]. These DIFs capture transferable representations that are less sensitive to geological variability, providing a more robust basis for downstream prediction. Prior work has demonstrated the utility of this strategy in engineering applications. For instance, Hu et al. [41] improved cross-domain performance by integrating DANN into ResPointNet++ to align training and test distributions. Ye et al. [42] demonstrated improved adaptability to varying drilling conditions using a DANN-based predictor. By learning highly transferable DIFs, DANN offers a principled route to strengthening generalization in complex, cross-project excavation scenarios.

In summary, this study proposes PL-DAST, a physics-constrained and LLM-driven domain-adversarial spatiotemporal framework for deformation prediction in deep foundation pits. PL-DAST incorporates three key innovations. First, it embeds geotechnical mechanisms into training via a physics-constrained loss function, enforcing physically plausible predictions. Second, it uses an LLM to parse and encode unstructured engineering text, enabling effective fusion of multi-source information. Third, it incorporates domain-adversarial training through a DANN to learn DIFs, improving cross-project transferability under varying geological conditions. Accordingly, the research addresses three key challenges: (1) How can geotechnical principles be embedded within predictive models to ensure physical consistency; (2) How can textual information, such as geological reports, be transformed into informative semantic embeddings using LLMs; and (3) How can DIFs be generated to ensure robust performance across diverse geological contexts. Applied to real-world cases, PL-DAST delivers deformation predictions that are both accurate and physically credible, supporting its use in excavation safety monitoring and risk-informed decision-making.

The remainder of this paper is structured as follows. Section 2 introduces the proposed PL-DAST framework and its implementation. Section 3 validates PL-DAST using three deep foundation pit projects from the Shanghai Metro. Section 4 performs comparative experiments to quantify the advantages of PL-DAST. Section 5 summarizes the key findings and outlines directions for future research.

2. Methodology

This section provides a comprehensive description of the PL-DAST framework, detailing the LLM-enhanced domain-adversarial feature learning process, the physics-constrained spatiotemporal prediction model, and the progressive three-stage training strategy designed to balance data-driven accuracy with geotechnical consistency.

2.1. Architecture of PL-DAST framework

This study proposes PL-DAST, a physics-constrained and large language model (LLM)-driven domain-adversarial spatiotemporal framework for accurate cross-project prediction of excavation-induced deformation. As illustrated in Fig. 1, PL-DAST adopts an encoder-decoder architecture that integrates textual information and physical constraints, which is structured as follows:

The encoder consists of an LLM-enhanced DANN to address the limited generalization of conventional DL models under cross-project distribution shifts. The DANN learns DIFs through adversarial training, improving transferability across heterogeneous excavation sites. Herein, the Qwen3 LLM encodes geological investigation reports into high-dimensional semantic embeddings, providing fine-grained contextual guidance for feature learning. A weighted residual connection then extracts features from the source and target domains separately and adaptively combines them to reduce distributional divergence. During adversarial training, the discriminator minimizes the discrimination loss (L_{disc}), while the feature extractor minimizes the adversarial loss (L_{adv}) together with the prediction accuracy loss (L_{acc}), encouraging representations that are both predictive and domain-invariant. Through this iterative process, the encoder progressively refines robust and generalizable feature representations, enabling more reliable prediction performance across diverse engineering contexts.

Building upon the DIFs extracted by the encoder, the decoder constructs a physics-constrained spatiotemporal prediction model for excavation-induced deformation. Specifically, it adopts an adaptive graph convolutional recurrent network (AGCRN) to generate end-to-end deformation prediction by modeling coupled spatiotemporal

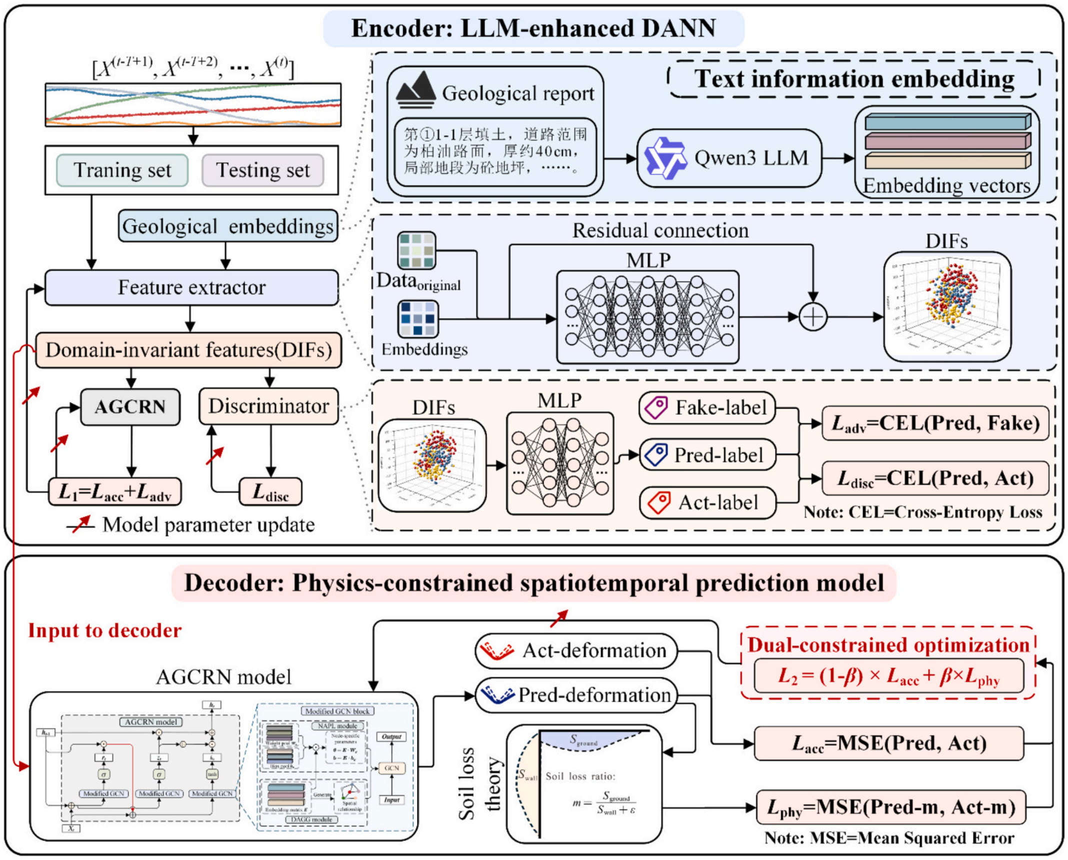

Fig. 1. Flowchart of the proposed PL-DAST framework.

dependencies in monitoring data. Training follows a dual-objective strategy that balances numerical fidelity with physical plausibility. The accuracy loss (L_{acc}) penalizes deviations between predicted and observed deformation values, while the physics-constrained loss (L_{phy}) constrains predictions through soil loss theory. Specifically, predicted deformation trajectories are input to a soil loss calculation module to derive the soil loss ratio m, a key physical indicator, which is then compared with its observed counterpart. The final objective function (L_2) combines L_{acc} and L_{phy} via a weighted formulation. By embedding geotechnical mechanisms directly into learning, the decoder produces deformation predictions that are both accurate and physically credible, supporting robust cross-project prediction in deep foundation pit engineering.

2.2. LLM-enhanced domain-adversarial feature learning

Although DL has shown great potential in underground engineering, its performance often degrades under cross-project distribution shifts, limiting practical deployment. To address this challenge, this study develops an LLM-enhanced DANN to extract DIFs through adversarial training. Specifically, the network consists of a feature extractor and a domain discriminator. The extractor generates feature representations with domain-invariant properties via multilayer nonlinear transformations, while the discriminator assesses distributional alignment by minimizing domain classification errors. Training follows an alternating optimization scheme in which the extractor is encouraged to confuse the discriminator, progressively suppressing project-specific information and reducing cross-project divergence. This adversarial learning process yields robust and transferable representations, thereby improving generalization across heterogeneous deep foundation pit projects.

To further enhance cross-project generalization, this study employs Qwen3 LLM to semantically parse geological investigation reports and convert them into high-dimensional embedding vectors that encode stratigraphic and geotechnical priors [43]. These semantic representations are then incorporated into the input of the feature extractor, thereby providing fine-grained geological context to guide domain-adversarial training. This multimodal fusion strengthens the representational capacity of the DANN under heterogeneous geological conditions, improving transferability while preserving data-driven accuracy and geological consistency. The development steps of the LLM-enhanced DANN are as follows:

2.2.1. LLM-based embedding of geological investigation reports

Geological investigation reports provide essential subsurface information, including soil types and stratigraphic thicknesses. Yet, conventional DL models are not designed to process unstructured text and therefore often reduce such records to coarse categorical inputs or ignore them entirely. This leads to a loss of valuable geological context and limits model transferability across projects with heterogeneous ground conditions. In response, LLMs offer an effective alternative. Pretrained on large-scale corpora, LLMs can capture contextual semantics and extract domain-relevant knowledge, including implicit geological attributes embedded in natural-language descriptions. By encoding investigation reports into high-dimensional semantic embeddings that preserve stratigraphic and geotechnical priors, LLMs translate unstructured narratives into informative numerical representations. These embeddings provide fine-grained geological context for domain-adversarial learning and downstream prediction, thereby improving robustness, interpretability and cross-project generalization in deep foundation pit deformation prediction.

Fig. 2 details the Qwen3 LLM-based process to transform unstructured geological investigation reports into machine-readable representations. Specifically, the Qwen3 embedding 4B LLM is deployed in a frozen inference-only mode, which leverages the extensive generalized

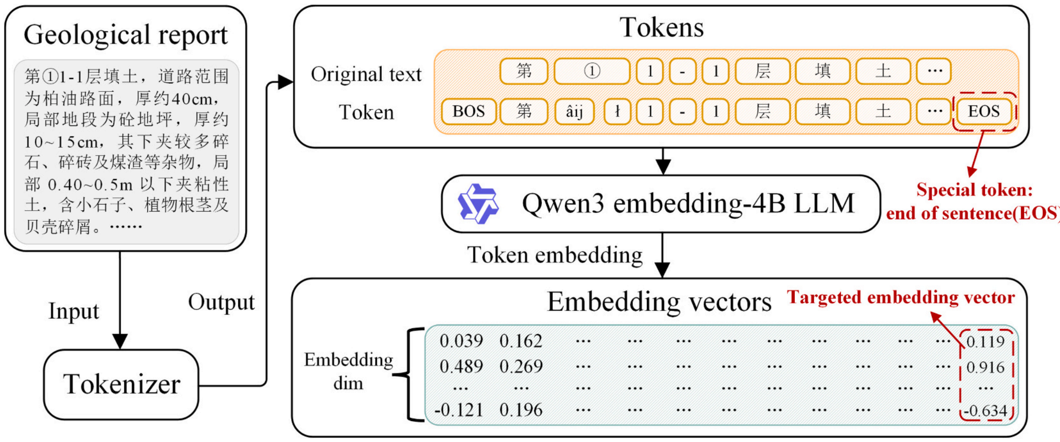

Fig. 2. Qwen3 LLM-based embedding process for geological investigation reports.

knowledge encoded within the pretrained model while maintaining computational efficiency during feature extraction. First, the raw text is tokenized into a sequence, with two special tokens BOS (Begin of Sentence) and EOS (End of Sentence) inserted to mark boundaries. Qwen3 then maps each token into a distributed semantic embedding, translating stratigraphic, lithological and geotechnical descriptions into high-dimensional numerical features. Notably, because Qwen3 is pre-trained to aggregate global contextual information at the EOS token, the corresponding hidden state is extracted to provide a compact representation of the full report.

Furthermore, given that the original embedding dimension of 4096 is substantial, this extracted vector is manually truncated to a 768-dimensional space. Concurrently, the sequential monitoring data is also projected into an equivalent 768-dimensional space. This targeted dimensionality alignment effectively decreases computational complexity and ensures a balanced representational capacity between the two modalities during the subsequent fusion phase, preventing the dense semantic priors from overwhelming the numerical constraints. By operating within this shared and balanced feature space, the processed LLM-derived embeddings establish a mathematically stable bridge between unstructured geotechnical semantics and quantitative numerical modeling. This structural alignment ultimately enables the effective fusion of geological textual evidence with dynamic monitoring data to support reliable cross-project domain adaptation.

2.2.2. DANN for learning domain-invariant features

The cross-project data distribution discrepancy constitutes a critical bottleneck limiting the generalization of existing DL methods, particularly in deep foundation pit engineering where geological conditions vary substantially between sites. DANNs address this limitation by learning DIFs through adversarial training. As illustrated in Fig. 3, the DANN architecture comprises three core components: (1) a feature extractor that maps inputs to DIFs, (2) a domain discriminator that predicts project identity to assess domain separability, and (3) an adversarial training mechanism that links the two, suppressing project-specific signals and promoting transferable representations.

During adversarial training, the feature extractor and domain discriminator engage in a minimax game. The discriminator is trained to identify the project origin of the extracted features, whereas the feature extractor is trained to confound this classification. This dynamic competition drives the extractor to suppress project-specific information and learn representations that are increasingly indistinguishable across domains, thereby reducing distributional discrepancies among heterogeneous datasets. As a result, DANN produces robust and transferable features that improve generalization in cross-project deep foundation pit prediction. The development steps of DANN are described below:

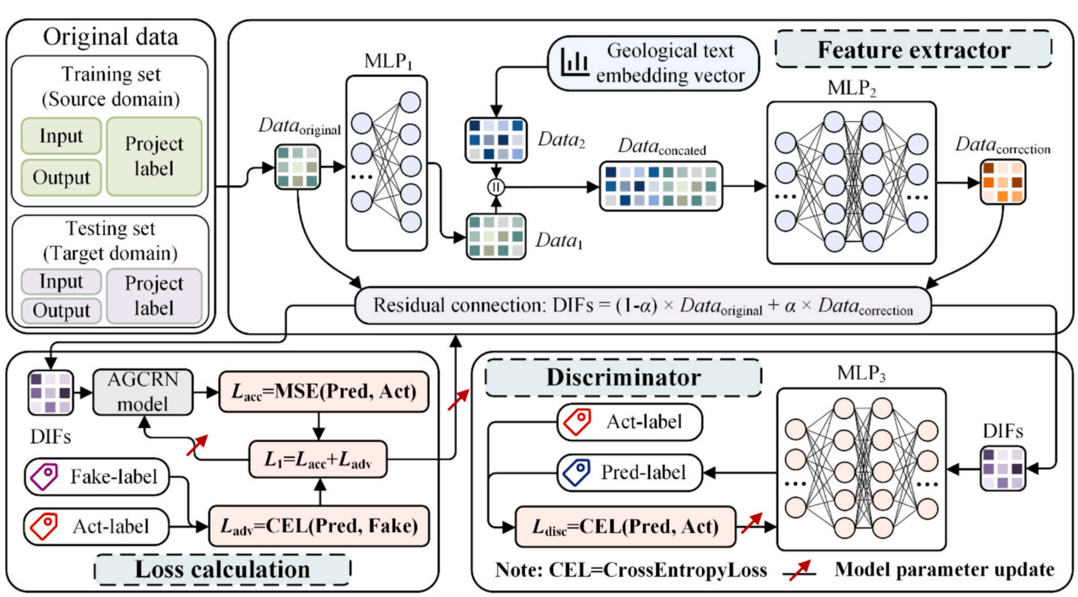

Fig. 3. Architecture of DANN for domain-invariant feature learning.

(1) Feature extractor.

The feature extractor is the core component of DANN and is designed to mitigate cross-project distribution shifts by learning DIFs through deep transformations. As shown in Fig. 3, the extractor follows a multi-stage processing pipeline. First, the original inputs ($\mathrm{Data_{original}}$) are mapped into a high-dimensional feature space using a multilayer perceptron ($\mathrm{MLP_1}$), yielding $\mathrm{Data_1}$. This representation is then concatenated with the LLM-derived geological embeddings $\mathrm{Data_2}$ and passed through a second multilayer perceptron $\mathrm{MLP_2}$ to generate a feature correction term ($\mathrm{Data_{correction}}$). To maintain a highly manageable model complexity, $\mathrm{MLP_2}$ is explicitly designed with compact hidden layers, employing ReLU activation functions in the intermediate layers to introduce essential non-linearity and a tanh activation function in the final layer, thereby optimizing feature extraction without incurring excessive computational overhead. This dual-path design enables joint exploitation of monitoring signals and text-derived geological knowledge. Finally, a weighted residual connection (Eq. (1)) adaptively fuses the original and corrected features, promoting cross-project alignment while retaining informative content. Overall, these components produce richer and more transferable representations that support robust downstream deformation prediction.

Parameters of the feature extractor are updated through a joint loss function, defined as $\mathrm{L_1} = \mathrm{L_{acc}} + \mathrm{L_{adv}}$ (Eq. (2)), where $\mathrm{L_{acc}}$ is the downstream prediction loss and $\mathrm{L_{adv}}$ is the adversarial loss. The accuracy term $\mathrm{L_{acc}}$ ensures that the learned DIFs remain informative for downstream prediction (see Section 2.3.3 for details). The adversarial term $\mathrm{L_{adv}}$ is computed using both real and fake labels (Eq. (3)), encouraging the extractor to produce representations that confuse the discriminator and thereby reduce cross-project distributional discrepancies. The fake labels are generated by randomly replacing the true labels to introduce sufficient divergence for effective alignment. The combined effect of two losses promotes features that are both predictive and transferable, thereby improving model stability and generalization across projects.

$$\mathrm{DIFs} = (1-\alpha)\times Data_{\mathrm{original}} + \alpha \times Data_{\mathrm{correction}} \tag{1}$$

where $Data_{\mathrm{original}}$ is the original input data; $Data_{\mathrm{correction}}$ is the correction terms generated by feature extractor.

$$L_1 = L_{\mathrm{acc}} + L_{adv} \tag{2}$$

$$L_{\mathrm{adv}} = \mathrm{CrossEntropyLoss}\left(label_{\mathrm{pred}}, label_{\mathrm{fake}}\right) \tag{3}$$

where L_1 is the total loss; L_{acc} is the accuracy loss; L_{adv} is the adversarial loss; $label_{\mathrm{pred}}$ are the project labels predicted by discriminator; and $label_{\mathrm{fake}}$ denotes randomly assigned fake labels that are deliberately different from the true project labels.

(2) Domain discriminator.

The domain discriminator is a key component of DANN and interacts adversarially with the feature extractor to promote the learning of DIFs. Implemented as a multilayer perceptron ($\mathrm{MLP_3}$), it takes DIFs as input, with an input dimension matching the extractor output, and produces one output node per engineering domain. To explicitly detail the network architecture, $\mathrm{MLP_3}$ employs ReLU activation functions across all hidden layers to capture complex nonlinear mappings, whereas its final layer intentionally omits any activation function to output raw logits directly. These logits are then evaluated by the cross-entropy loss function to compute the probability that a feature representation originates from each domain. The highest-probability domain is taken as the predicted label, providing a quantitative measure of domain separability and generating the adversarial signal that drives cross-project feature alignment.

The discriminator is trained exclusively by minimizing the discrimination loss $\mathrm{L_{disc}}$ (Eq. (4)), defined as the cross-entropy between predicted and true domain labels. This objective progressively improves its ability to distinguish domains. The discriminator's prediction accuracy, together with the feature extractor's deception rate (Eqs. (5)–(6)), provides a direct indication of the degree of domain invariance achieved by the learned DIFs and delivers clear optimization signals for the extractor. Through this feedback loop, the discriminator and feature extractor are updated in opposing directions, yielding a stable adversarial process that iteratively refines both modules and produces more robust and transferable representations.

$$L_{\mathrm{disc}} = \mathrm{CrossEntropyLoss}\left(label_{\mathrm{pred}}, label_{\mathrm{act}}\right) \tag{4}$$

where $label_{\mathrm{act}}$ are the actual project labels.

$$\mathrm{Accuracy} = \frac{1}{S}\sum_{i=1}^{S} I\left(label_{\mathrm{pred}}^{(i)} = label_{\mathrm{act}}^{(i)}\right) \tag{5}$$

$$\mathrm{Deception\ rate} = 1 - \mathrm{Accuracy} \tag{6}$$

where S is the total number of samples; $I(\cdot)$ is the indicator function, which equals 1 when the condition is true and 0 otherwise; $label_{\mathrm{pred}}^{(i)}$ and $label_{\mathrm{act}}^{(i)}$ are the predicted and true labels of the ith sample, respectively; Accuracy is the prediction accuracy of discriminator; Deception rate quantifies the probability that the feature extractor successfully deceives the discriminator by producing DIFs.

(3) Adversarial training.

Adversarial training between the feature extractor and domain discriminator constitutes the core mechanism of DANN. The two modules are optimized by independent loss functions: $\mathrm{L_1}$ for the feature extractor and $\mathrm{L_{disc}}$ for the discriminator. This decoupled design assigns complementary roles, with the extractor learning DIFs and the discriminator learning to distinguish project domains. Training follows an alternating optimization: the discriminator is updated with the extractor fixed, and the extractor is then updated while keeping the discriminator parameters frozen.

This adversarial process converges toward a dynamic equilibrium. As the discriminator improves its ability to distinguish domains, it reveals residual project-specific signals in the learned representations, prompting the feature extractor to suppress these cues more strongly. Through repeated adversarial updates, the extractor progressively learns features with increasing domain invariance, leading to closer alignment of cross-domain distributions. Ultimately, this optimization process enables DANN to acquire highly transferable representations and substantially enhances generalization across engineering applications.

2.3. Physics-constrained spatiotemporal deformation prediction model

Building upon the DIFs generation method introduced in Section 2.2, the next goal is to make accurate spatiotemporal deformation prediction in deep foundation pits. Conventional DL approaches exhibit limited generalization for two main reasons. First, monitoring data distributions vary substantially across projects, reducing transferability under cross-site deployment. Second, most models are trained purely from data and do not explicitly encode geotechnical physics, increasing the likelihood of predictions that deviate from physically plausible behavior.

To address these limitations, this study proposes a physics-constrained spatiotemporal prediction model that embeds soil loss theory as an explicit constraint in the training objective. This formulation enforces adherence to geotechnical principles while retaining the flexibility of data-driven learning. In parallel, the model leverages the inherent spatiotemporal dependencies in deformation monitoring data to improve prediction accuracy. By coupling physical priors with learned data patterns, the proposed model enhances interpretability and improves prediction reliability and generalization in cross-project deep foundation pit prediction. The physics-constrained spatiotemporal prediction is described below.

2.3.1. Soil loss theory as a physical constraint

As a physically grounded constraint, the soil loss theory provides a mechanistic basis for ensuring that data-driven predictions remain consistent with the spatial redistribution of ground movements during excavation. Originally proposed by Liu and Hou [44], it is a semi-empirical approach for estimating deformation in deep excavations. As shown in Fig.4, its theoretical foundation rests on a correlation formula in Eq. (7) that links retaining wall deformation to ground surface settlement through the principle of soil movement area. This formulation captures key soil-structure interaction behavior in a tractable form, enabling rapid estimation of deformation profiles without requiring complex constitutive modeling.

While the method has proven valuable in many contexts, certain aspects may benefit from further refinement in complex or data-scarce scenarios. First, wall deformation and settlement curves are typically computed using finite rod elements or the elastic foundation beam method, which rely on simplified mechanics and do not exploit monitoring data. Second, the soil loss ratio m, a key parameter, is commonly prescribed as an empirical constant, limiting its ability to represent the time-dependent evolution of excavation-induced deformation.

To broaden the applicability of soil loss theory to realistic excavation scenarios, this study proposes a modified formulation that enhances both adaptability and physical relevance. Conventional mechanics-based calculations are replaced with data-driven estimators, reducing discrepancies between theoretical predictions and field observations. In addition, the soil loss ratio m is reformulated as a time-varying parameter and updated dynamically via time-series forecasting to capture excavation-induced disturbances. These modifications transform soil loss theory from a static empirical tool into a physically consistent and adaptive constraint. This enables its seamless integration into physics-constrained learning, where it enforces kinematic compatibility of the deformation field and improves the physical plausibility of data-driven predictions for deep foundation pit deformation.

$$m \times (S_{\text{wall}} + \varepsilon) = S_{\text{ground}} \tag{7}$$

where m is the soil loss ratio; S_{wall} is the ground movement area associated with retaining wall deformation; S_{ground} is the ground movement area associated with ground surface settlement; and ε is a small constant (set to 1) introduced to stabilize optimization and mitigate numerical instability during the early stages of neural network training.

2.3.2. AGCRN-based spatiotemporal deformation prediction

To fully leverage the spatiotemporal characteristics of deformation monitoring data, this study employs the adaptive graph convolutional recurrent network (AGCRN) proposed by Bai et al. [45] as the core predictor. As shown in Fig. 5, AGCRN extends the gated recurrent unit (GRU) formulation (Eqs. (8)–(11)) by replacing the fully connected transformations with a modified graph convolution module. This design enables simultaneous learning of spatial interactions and temporal dynamics. Notably, AGCRN can learn latent spatiotemporal dependencies in multivariate time series without requiring a predefined graph, supporting more accurate and robust deformation prediction in deep foundation pits.

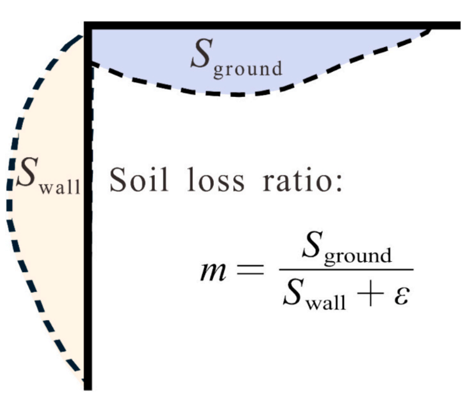

Fig. 4. Schematic illustration of the soil loss theory linking retaining wall deformation and ground surface settlement.

The modified GCN in AGCRN consists of two core modules: node-adaptive parameter learning (NAPL) and data-adaptive graph generation (DAGG). NAPL learns node-specific convolutional parameters from the node embedding matrix E using a weight pool W_g, and a bias pool b_g (Eqs. (12)–(13)), enabling the model to capture heterogeneity across monitoring points and improving convolution fidelity. In parallel, DAGG constructs the adjacency matrix $\widetilde{A}$ adaptively from the embedding matrix E (Eq. (14)), allowing latent spatial correlations to be identified without assuming a fixed graph structure. These outputs are then integrated through the graph convolution operation (Eq. (15)) to yield unified spatiotemporal representations.

In short, the dual-module design enables AGCRN to learn node-specific convolutional parameters via NAPL while constructing adaptive adjacency structures through DAGG. This combination captures local heterogeneity and reveals latent spatial dependencies that are often missed by conventional graph-based models with fixed connectivity. As a result, AGCRN provides more accurate, robust and generalizable predictions, making it well suited for modeling complex spatiotemporal deformation in deep foundation pits.

$$r_t = \sigma[\text{Modified GCN}(W_r \cdot (X_t \cdot h_{t-1}) \oplus b_r)] \tag{8}$$

$$z_t = \sigma[\text{Modified GCN}(W_z \cdot (X_t \cdot h_{t-1}) \oplus b_z)] \tag{9}$$

$$\widetilde{h}_t = \tanh[W_h(X_t \cdot (r_t \odot h_{t-1}) \oplus b_h)] \tag{10}$$

$$h_t = ((1 - z_t) \odot h_{t-1}) \oplus \left(z_t \odot \widetilde{h}_t\right) \tag{11}$$

where X_t denotes the input data at time step t; W_r, W_z, and W_h are the weight matrices for the reset gate, update gate, and candidate hidden state, respectively; b_r, b_z, and b_h are the bias vectors for the reset gate, update gate, and candidate hidden state, respectively; r_t and z_t are the reset gate and update gate at time step t, respectively; σ denotes the sigmoid activation function; tanh denotes the hyperbolic tangent activation; $\widetilde{h}_t$ is the candidate hidden state at time step t; h_t is the hidden state at time step t.

$$W = E \cdot W_g \tag{12}$$

$$b = E \cdot b_g \tag{13}$$

where E is the node embedding matrix; W_g and b_g are the weight and bias pool, respectively; W and b are node-specific convolutional parameters (weights and biases) generated by the NAPL module.

$$\widetilde{A} = I_N + \text{Softmax}\left(\text{ReLU}\left(E \cdot E^T\right)\right) \tag{14}$$

where $\widetilde{A}$ is the adjacency matrix generated by the DAGG module; I_N is an $N \times N$ identity matrix.

$$Output = \sigma\left(\widetilde{A} \cdot Input \cdot W \oplus b\right) \tag{15}$$

2.3.3. Physics-constrained training strategy and loss design

Building upon the soil loss theory and the AGCRN-based spatiotemporal predictor, this study introduces a physics-constrained training strategy. The motivation for this design stems from the limitations of purely data-driven models, which may yield numerically accurate results yet violate geotechnical constraints and generalize poorly to unseen projects. To address this, our central contribution is to embed soil loss theory as an explicit physical constraint within the AGCRN training objective. This explicit physics constraint steers the training process to ensure that the predicted deformations remain physically reasonable while preserving the adaptability of DL. As a result, the model attains

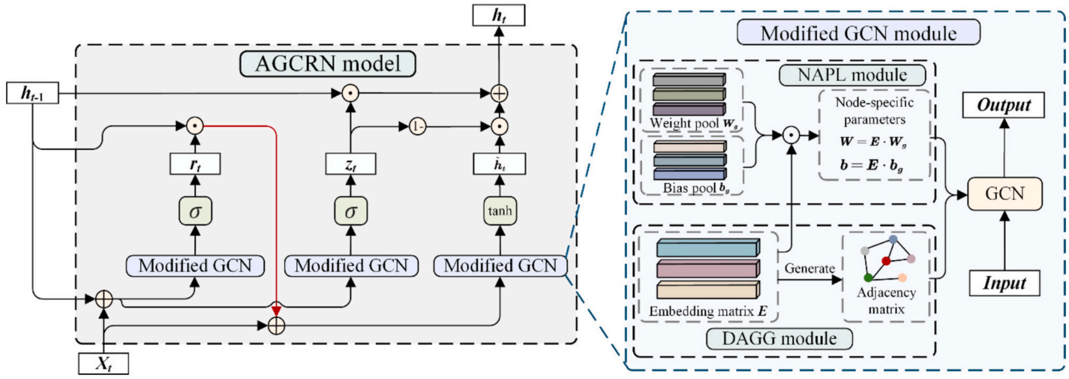

Fig. 5. AGCRN architecture for spatiotemporal deformation prediction.

improved reliability, interpretability and cross-project generalization in different deep foundation pit scenarios. The implementation comprises three key components:

(1) Accuracy loss (L_{acc}) calculation.

To quantify prediction performance for two key deformations induced by deep excavation (retaining wall deformation and ground surface settlement), this study defined L_{acc} in Eq. (16). This loss function employs the mean square error (MSE) to systematically measure deviations between predicted values and monitoring observations, thereby providing a direct data-driven signal for parameter optimization.

(2) Physics-constrained loss (L_{phy}) calculation.

The soil loss ratio m is the key parameter in soil loss theory and encodes the relationship between retaining wall deformation and ground surface settlement. To impose physical constraints during training, this study defines a physics-constrained loss L_{phy} based on the prediction error of m. Importantly, m is not predicted directly by the network. Instead, it is derived from the predicted deformation fields, ensuring that the constraint acts on the deformation trajectories themselves. The calculation proceeds in two steps. First, numerical integration is applied to the predicted deformation curves to obtain the corresponding ground movement areas. Specifically, given that the monitoring data consists of discrete values, the trapezoidal rule is employed to approximate the definite integrals. For the retaining wall, the deformation area S_{wall} is calculated by integrating the horizontal displacement values with respect to the depth axis (y-axis). Similarly, for the ground surface settlement, the area S_{ground} is computed by integrating the vertical settlement values along the horizontal distance axis (x-axis). This approach allows discrete engineering monitoring data to be directly incorporated into the physical constraint, ensuring that the calculated soil loss ratio m accurately reflects the real-time volume of soil displacement. Second, the predicted soil loss ratio m is calculated using Eq. (7). The physics-constrained loss L_{phy} in Eq. (17) is then calculated by comparing this predicted value with the corresponding observed value of m, providing an explicit physics-based signal for parameter optimization.

(3) Accuracy-physics hybrid loss (L_2) calculation.

The composite loss L_2 (Eq. (18)) is defined as a weighted linear combination of the accuracy loss L_{acc} and the physics-constrained loss L_{phy}, enabling joint optimization of data fidelity and physical plausibility. Here, L_{acc} promotes numerical agreement with monitoring data, while L_{phy} constrains predictions to satisfy soil loss theory. This physics-constrained training strategy ensures that the soil movement areas associated with retaining wall deformation and ground surface settlement remain physically reasonable, resulting in predictions that are both numerically accurate and engineering-credible.

$$L_{acc} = \mathrm{MSE}(\hat{y}, y) = \frac{1}{n}\sum\nolimits_{i=1}^{n}(y_i - \hat{y}_i)^2 \tag{16}$$

$$L_{phy} = \mathrm{MSE}(\hat{m}, m) = \frac{1}{n}\sum\nolimits_{i=1}^{n}(m_i - \hat{m}_i)^2 \tag{17}$$

$$L_2 = (1-\beta)L_{acc} + \beta L_{phy} \tag{18}$$

where y_i and $\hat{y}_i$ represent the observed and predicted deformation values of the ith sample, respectively; m_i and $\hat{m}_i$ represent the observed and predicted soil loss ratio for the ith sample, respectively; n denotes the total number of samples; β is the weight coefficient for L_{phy}, which is set to 0.3 in this study.

2.4. Training and evaluation of PL-DAST

This section outlines the comprehensive training procedure and the corresponding evaluation methodologies for the proposed framework. Specifically, it first introduces a progressive three-stage training strategy designed to stabilize optimization and effectively balance data-driven learning with geotechnical mechanisms. Subsequently, it details the evaluation metrics used to quantify both cross-project feature alignment and spatiotemporal deformation prediction accuracy.

2.4.1. Three-stage training strategy

Building on the methods in Sections 2.2 and 2.3, this study develops a PL-DAST framework that integrates domain-invariant feature learning with physics-constrained geotechnical constraints. To achieve stable convergence and realize complementary gains in accuracy, transferability and physical consistency, a three-stage training strategy is designed in Fig. 6, as detailed below.

(1) Stage 1 (epochs 1–20): Collaborative training of DANN and AGCRN.

This stage adopts a dual-level optimization strategy. At the first level, the feature extractor and AGCRN are jointly trained using a composite loss $L_1 = L_{acc} + L_{adv}$, where L_{acc} promotes predictive accuracy and provides an initial supervisory signal for AGCRN, and L_{adv} improves their domain invariance in the learned representations. At the second level, the discriminator is trained independently by minimizing L_{disc} to strengthen domain classification capability. This dual-level optimization establishes a domain-invariant feature space while preserving predictive accuracy, providing a solid foundation for subsequent integration of physics constraints.

(2) Stage 2 (epochs 21–30): Accuracy-oriented pretraining of AGCRN.

Building on the feature alignment achieved in Stage 1, this stage freezes the feature extractor parameters to focus exclusively on AGCRN model training. To avoid instability from introducing physical constraints prematurely, training in this stage minimizes L_{acc} alone, yielding a purely data-driven fit. This accuracy-oriented pretraining guides the model toward a well-conditioned solution space, reduces the risk of collapse during subsequent physics-constrained optimization, and

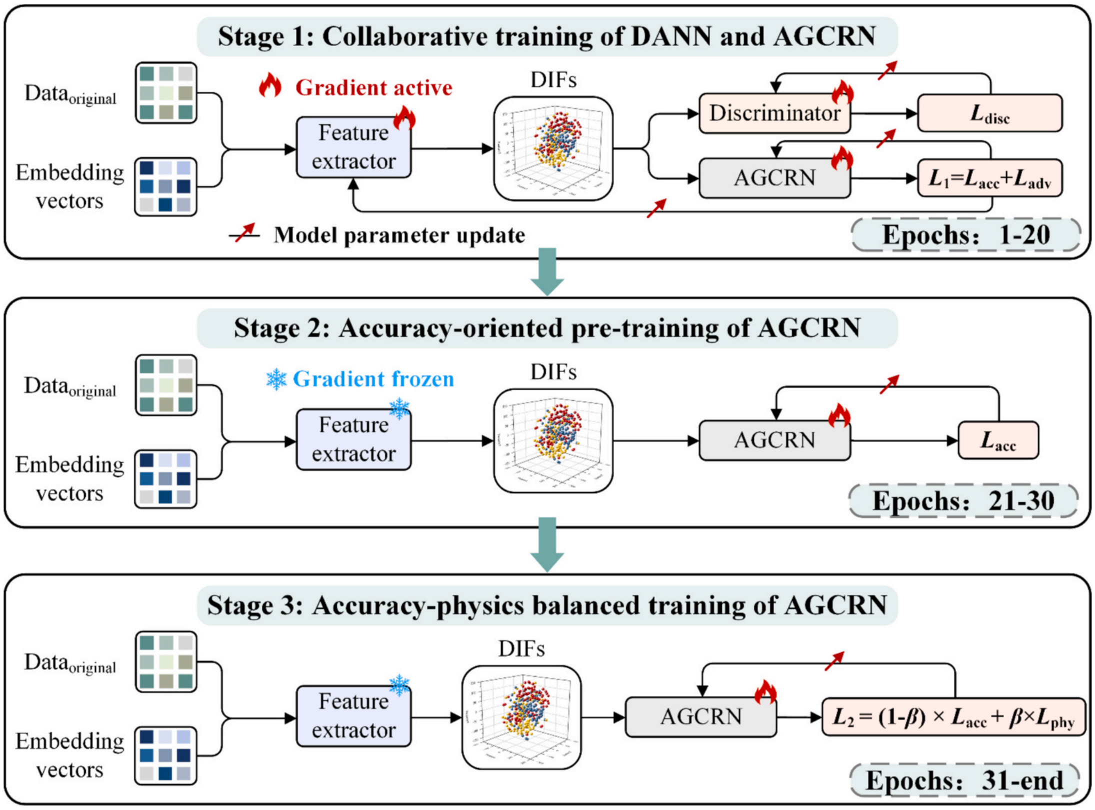

Fig. 6. Three-stage training strategy for the PL-DAST framework.

promotes stable convergence.

(3) Stage 3 (epoch 31 to convergence): Accuracy-physics balanced training of AGCRN.

The accuracy-oriented solution obtained in Stage 2 may overfit project-specific patterns and thus limit cross-project generalization. To mitigate it, this stage introduces a hybrid loss L_2 that balances data fidelity with physical consistency. By embedding soil loss theory as a geotechnical constraint, training is guided beyond purely data-driven local optima, yielding deformation predictions that remain both accurate and physically plausible across different engineering scenarios.

Overall, the staged training strategy, including collaborative training, accuracy-oriented pretraining, and accuracy-physics balancing, stabilizes physics-constrained learning, reduces the risk of training collapse, and guides convergence toward robust and high-fidelity solutions. By progressively integrating data-driven learning with geotechnical constraints, the PL-DAST framework can deliver spatiotemporal deformation predictions that are both reliable and engineering-credible across diverse excavation scenarios.

2.4.2. Evaluation metrics

Training performance is monitored after each epoch by evaluating predictive accuracy, with a coefficient of determination R^2 exceeding 90%. The reported metrics include the discriminator's classification accuracy, the feature extractor's deception rate, and the loss terms L_2, L_{acc}, and L_{phy}, as defined above. To quantify cross-project feature alignment, the Sliced-Wasserstein Distance (SWD) is employed. Direct computation of the Wasserstein distance becomes prohibitively expensive in high-dimensional feature spaces, as the computational cost increases rapidly with dimensionality. SWD offers an efficient approximation by projecting the feature distributions onto 200 randomly sampled 1-dimensional subspaces and averaging the resulting Wasserstein distances. This approach substantially reduces computational cost while remaining sensitive to distributional discrepancies. The SWD between two feature sets is defined in Eq. (19) and is widely used in domain adaptation and feature alignment tasks.

While SWD and the loss terms quantify feature alignment and training dynamics at the representation level, the ultimate goal of PL-DAST is to achieve accurate spatiotemporal deformation prediction. To this end, two task-specific metrics are employed. The time-averaged absolute error (TAE) measures the mean absolute error over time at each spatial location (Eq. (20)), capturing temporal reliability. The spatially averaged MAE (SA-MAE) computes the mean absolute error across all spatial points at each time step (Eq. (21)), assessing spatial consistency. These metrics provide a balanced and comprehensive assessment of spatiotemporal prediction robustness.

$$\mathrm{SWD}(X,Y)=\frac{1}{L}\sum_{l=1}^{L}W_1\left(P_{\theta_l}X,P_{\theta_l}Y\right) \tag{19}$$

where $X=\{x_i\}_{i=1}^{n}$ and $Y=\left\{y_j\right\}_{j=1}^{m}$ denote the feature samples from two different domains, with $x_i,y_j\in\mathbb{R}^d$; $P_{\theta_l}(\cdot)=<\cdot,\theta_l>$ denotes the projection of samples onto the direction $\theta_\ell\in\mathbb{S}^{d-1}$; $W_1(\cdot,\cdot)$ is the 1-dimensional Wasserstein distance, computed as the mean absolute difference between sorted projected samples; L is the number of random projection directions, which is set to 200 in this study.

$$\mathrm{TAE}=\frac{1}{T}\sum_{i=1}^{T}|c_i-\widehat{c}_i| \tag{20}$$

$$\mathrm{SA-MAE}=\frac{1}{U}\sum_{j=1}^{U}|d_j-\widehat{d}_j| \tag{21}$$

where, T is the total number of days; U is the total number of measurement points; c_i and $\widehat{c}_i$ are the actual and predicted deformation at the given location on day i, respectively; d_j and $\widehat{d}_j$ are the actual and predicted deformations, respectively, at the ith location on the given day.

3. Case study

This section validates the effectiveness of the proposed PL-DAST framework using three real-world deep foundation pit projects from the Shanghai Metro. The experimental setup involves a rigorous cross-project evaluation protocol to assess the predictive accuracy and physical consistency of the model under heterogeneous geological conditions.

3.1. Project overview and data description

This subsection provides a detailed description of the engineering backgrounds and the data preparation process. It covers the site characteristics of the selected metro stations, the structure of the monitoring indicator system, and the engineering-agnostic preprocessing strategy used to align heterogeneous monitoring data into a consistent representational space.

3.1.1. Site description and excavation characteristics

The proposed PL-DAST framework is applied to three deep foundation pit projects in the Shanghai Metro: Guilin Park Station on Line 15, East China University of Science and Technology (ECUST) Station on Line 15, and a section of Lantian Road Station on Line 9, Shanghai, China. Site layouts are shown in Fig. 7, with excavation lengths of 160 m, 170 m, and 250 m, and widths of 24 m, 23 m, and 22 m, respectively. Multiple monitoring cross-sections are set at each site to satisfy risk-assessment requirements. All three projects are constructed using the open-cut excavation method, with excavation depths of 24.3 m, 16.0 m, and 16.5 m, respectively, and are therefore classified as deep foundation pits. Representative cross-sections and the corresponding support systems are shown in Fig. 8.

Retaining wall deformation and ground surface settlement are the primary deformation modes governing safety and serviceability in deep foundation pits. To mitigate these risks, all three construction sites employ a composite support system (Table 1), comprising retaining walls with thicknesses of 1200 mm, 800 mm, and 800 mm, respectively, in combination with hybrid concrete-steel supports. This rigid-flexible configuration satisfies structural safety requirements while remaining cost-effective. Notably, all sites are situated in densely built urban areas with sensitive surroundings, further highlighting the need for robust deformation control and reliable prediction.

3.1.2. Data description and preprocessing

Table 2 summarizes the indicator system used in the PL-DAST framework. Specifically, inputs consist of two monitored deformations responses: ground surface settlement at 100 points and retaining wall deformation at 101 points. The output includes the predicted values of these two deformation types, along with the soil loss ratio m. Specifically, PL-DAST uses observations from days $(t-4)$ to $(t-1)$ to predict the deformation response on day t. Notably, the soil loss ratio m is not predicted directly. Instead, it is computed by applying soil loss theory to the predicted deformation outputs from PL-DAST. This deliberate design avoids direct regression of m, ensuring that physical laws governing the soil loss relationship constrains deformation predictions in a mechanistically meaningful way, thereby strengthening interpretability and physical consistency.

In practical engineering projects, the number and layout of monitoring points vary across projects because excavation scale, instrumentation design and geological conditions differ. This variation prevents direct alignment of input dimensions and complicates the development of a unified cross-project predictor. Moreover, irregular sensor spacing introduces inconsistencies in spatial reference, which can obscure stable deformation patterns. To address these issues, PL-DAST employs an "engineering-agnostic" data preprocessing strategy. Using retaining wall deformation as an example (Fig. 9), the procedure first fits a continuous deformation profile to the measured data. It then resamples 101 uniformly spaced virtual points along a normalized depth axis and interpolates the corresponding deformation values from the fitted curve. This transformation maps heterogeneous observations into a dimensionally consistent and spatially aligned feature space, reducing project-specific artefacts while preserving the physical structure of the deformation field. Consequently, the indicator system improves structural

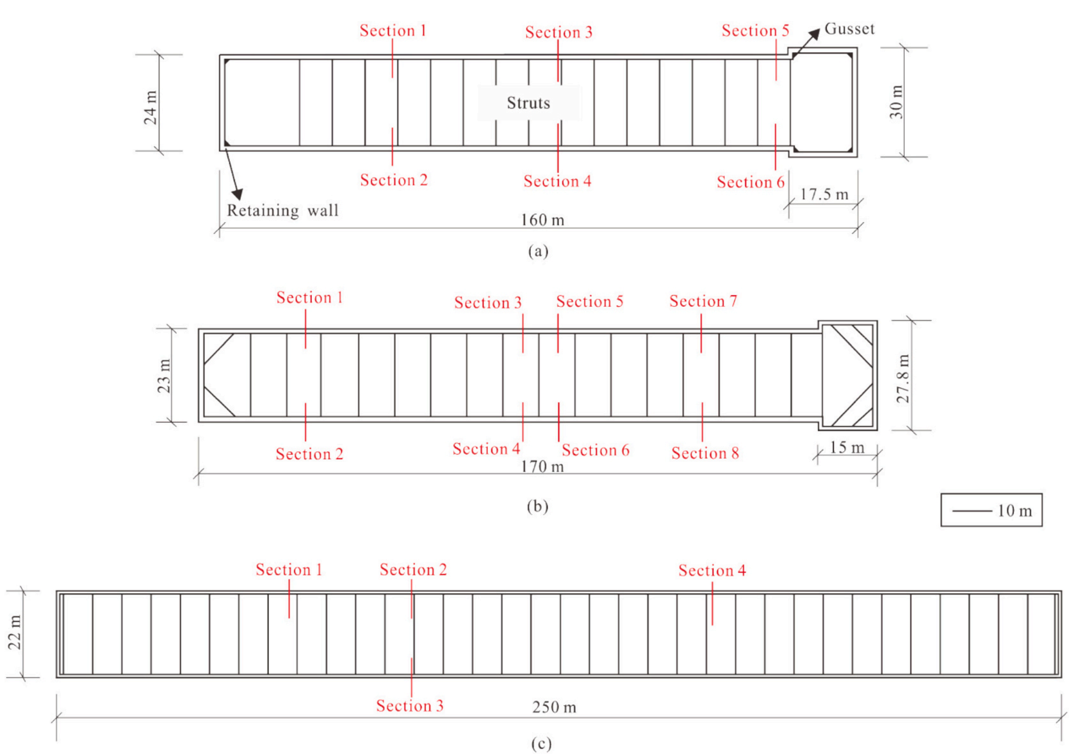

Fig. 7. Plan layouts of the three deep foundation pits: (a) Guilin Park Station (Line 15); (b) ECUST Station (Line 15); (c) Lantian Road Station (Line 9).

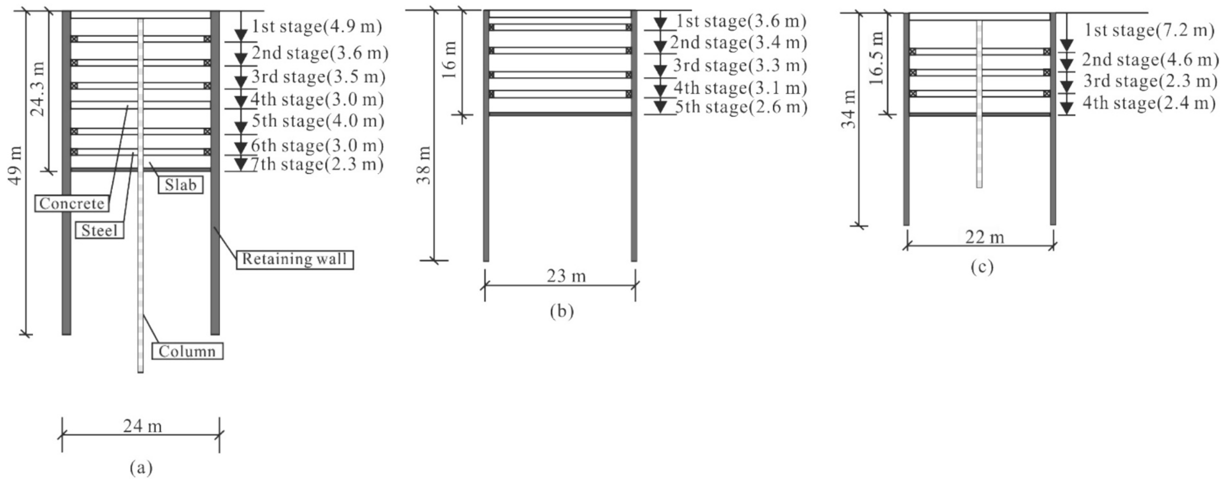

Fig. 8. Representative cross-sections and corresponding excavation stages for the three deep foundation pits: (a) Guilin Park Station (Line 15); (b) ECUST Station (Line 15); (c) Lantian Road Station (Line 9).

Table 1
Structural parameters of the support systems in the three deep foundation pits.

Foundation pit	Structure	Material	Size (mm)	Spacing (mm)	E (GPa)
Guilin Park Station	Diaphragm wall	Reinforced concrete	1200 mm	/	31
	First support	Concrete	1400 mm × 900 mm	8000	31
	Second - Fourth support	Steel	Φ609, t = 16 mm	2500	210
	Fifth support	Concrete	1400 mm × 900 mm	8000	31
	Sixth - Seventh support	Steel	Φ800, t = 20 mm	2500	210
ECUST Station	Diaphragm wall	Reinforced concrete	800 mm	/	31
	First support	Concrete	760 mm × 800 mm	9000	31
	Second - Fifth support	Steel	Φ609, t = 16 mm	3000	210
Lantian Road Station	Diaphragm wall	Reinforced concrete	800 mm	/	31
	First support	Concrete	800 mm × 900 mm	7300	31
	Second - Fourth support	Steel	Φ609, t = 16 mm	2500	210

Table 2
Overview of the indicator system and descriptive statistics used in the PL-DAST framework.

Feature	Data type	Dimension	Unit	Mean (±Std)	[Min, Max]
Ground surface settlement	Input & Output	100	mm	−10.11 (±8.89)	[−52.13, −0.01]
Retaining wall deformation	Input & Output	101	mm	13.71 (±13.52)	[−5.74, 73.00]
Soil loss ratio m	Output only	1	/	0.92 (±0.32)	[0.34, 1.73]

adaptability and supports generalization across diverse deep foundation pit projects.

Overall, this study adopts a strict cross-project evaluation protocol. Specifically, 781 samples from the Guilin Park and ECUST stations are used for training, whereas 111 samples from the Lantian Road station are held out as an independent test set. Although the test set constitutes only around 12.4% of the total samples, it originates from a different project and therefore provides a realistic proxy for deployment under unseen site conditions. This partitioning enables a rigorous assessment of model transferability and cross-project generalization.

3.2. PL-DAST framework implementation

This subsection details the concrete implementation steps of the PL-DAST framework across the study sites. The discussion focuses on the practical extraction of geological semantic embeddings using the Qwen3 LLM and the execution of the three-stage training strategy to achieve stable convergence for the physics-constrained model.

3.2.1. LLM-based extraction of geological semantic embeddings

Drawing on engineering geological investigation reports, this study employs the Qwen3 embedding-4B LLM to extract semantic information that is typically inaccessible to conventional models. Qwen3 encodes unstructured geological descriptions into 2560-dimensional embedding vectors that capture stratigraphic and geotechnical information. To balance semantic richness with computational efficiency, the first 768 dimensions of the EOS token embedding are retained as the final geological feature vector. This dimensional reduction preserves critical geological semantics while ensuring compatibility with the DANN feature extractors and improving training efficiency. By converting narrative reports into machine-readable representations, the LLM provides an effective bridge between textual records and numerical deformation prediction.

To quantify the contribution of LLM-derived text features to cross-project adaptability, Table 3 reports the cosine similarity of geological embedding vectors across the three projects. All similarity scores exceed 0.895, indicating that the embeddings capture stratigraphic signatures common to Shanghai sites. The values are nonetheless distinct, showing that the representations also retain project-specific information, such as differences in layer thickness and stratigraphic sequence. This combination of regional consistency and local sensitivity supports the use of LLM embeddings as a robust foundation for cross-project transfer learning.

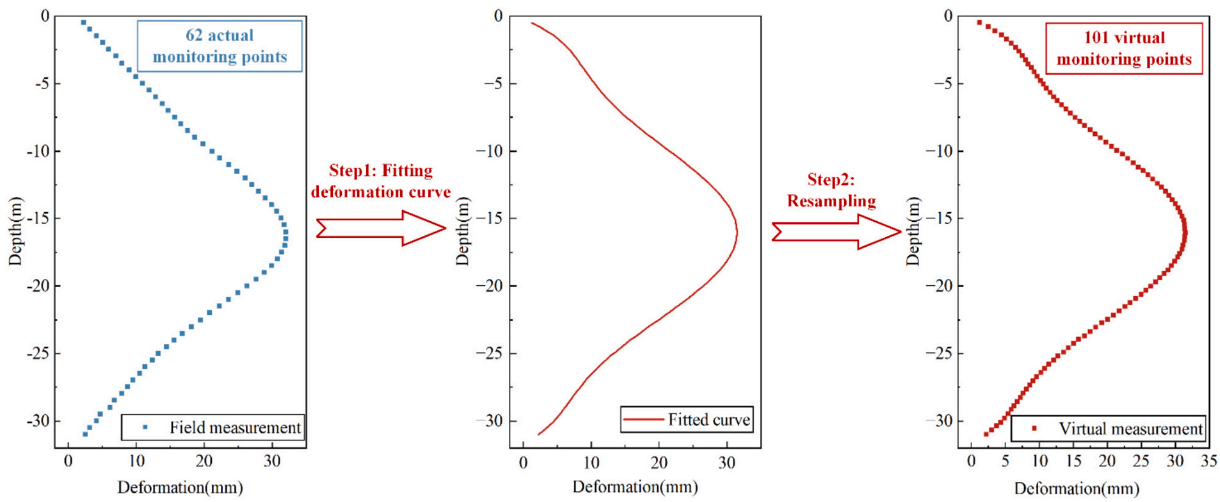

Fig. 9. Engineering-agnostic preprocessing of retaining wall deformation data.

Table 3
Cosine similarity of LLM-derived geological embeddings across three projects.

	Guilin Park Station	ECUST Station	Lantian Road Station
Guilin Park Station	1.000	0.920	0.895
ECUST Station	0.920	1.000	0.911
Lantian Road Station	0.895	0.911	1.000

3.2.2. *Three-stage training strategy and its advantages*

PL-DAST is optimized using the three-stage training strategy outlined in Section 2.4.1. Stage 1 jointly trains the LLM-enhanced DANN and AGCRN to learn domain-aligned representations. Stage 2 then refines the predictor through accuracy-oriented pretraining. Stage 3 introduces soil loss theory as the physical constraint to enforce physical consistency while maintaining predictive performance. This progressive optimization stabilizes training and guides convergence toward solutions that balance numerical accuracy with geotechnical plausibility. Hyperparameters of PL-DAST are tuned via Bayesian optimization over 100 iterations, and the optimal configuration are summarized in Table 4. To further detail the network architecture, the Adam optimizer is employed across all training stages. Regarding model complexity, the total number of trainable parameters in the PL-DAST framework is strictly lightweight at 78,439. All experiments are conducted on a high-performance computing platform (Intel W-2255 CPU, 64 GB RAM, RTX 3080Ti GPU). The full training procedure completes in 86 s, and inference requires only 0.01 s per prediction, indicating that the proposed physics-constrained and LLM-driven framework is computationally efficient for practical deployment. The evolution of key training indicators across the three-stage strategy is shown in Fig. 10.

Table 4
Hyperparameter search space and optimal configuration for the PL-DAST framework.

Hyperparameter	Module	Search space	Optimal combination
Hidden layer dimensions	AGCRN model	[16, 128]	24
Network layers		[2,4]	2
Chebyshev polynomial order		[1,2]	1
Embedding dimension of node embedding matrix *E*		[4,16]	14
Learning rate		[1e-3, 1e-1]	0.00290
Network layers	Feature extractor	[4,6]	6
Learning rate		[1e-4, 1e-3]	0.00020
Weight of $\text{Data}_{\text{correction}}$		[0.1, 0.3]	0.11868
Hidden layer dimensions		/	1536
Network layers	Discriminator	[2,3]	3
Learning rate		[1e-5, 1e-4]	0.00004
Hidden layer dimensions		[16, 64]	23
Epoch	Global	[40, 60]	60

As shown in Fig. 10(a), Stage 1 is characterized by a progressive decrease in discriminator accuracy and a concurrent increase in the feature extractor's deception rate. This inverse trend reflects effective adversarial learning: as the feature extractor learns to produce high-quality DIFs, the discriminator becomes increasingly unable to distinguish between source and target project domains. Consequently, the learned representations exhibit improved cross-project transferability, establishing a robust feature basis for subsequent spatiotemporal deformation prediction.

As shown in Fig. 10(b), the total loss exhibits pronounced fluctuations during Stage 1, reflecting the inherent instability of joint domain adaptation between the feature extractor and the AGCRN predictor. In Stage 2, the total loss decreases rapidly and converges toward a local minimum, indicating a more stable and accuracy-dominated optimization regime. Stage 3 further reduces the total loss, showing that introducing physical constraints improves training robustness while maintaining predictive fidelity.

As shown in Fig. 10(c), the accuracy loss follows the trend of the total loss but declines more steeply. Stage 1 exhibits high and unstable values, consistent with the uncertainty introduced by early cross-project domain adaptation. In Stage 2, the accuracy loss drops rapidly as supervised training dominates, driving the predictor toward an accuracy-oriented optimum. Stage 3 introduces modest oscillations when soil loss theory-based physical constraints are activated. Nevertheless, the accuracy loss remains low and stable, indicating that physics-constrained regularization enhances physical consistency without compromising numerical performance.

As shown in Fig. 10(d), the physics-constrained loss remains relatively high during Stages 1 and 2, indicating that accuracy-dominated optimization can lead to noticeable departures from the soil loss constraint. In Stage 3, once the physical constraint is activated, L_{phy} increases briefly as the model begins to incorporate the constraint, and then decreases steadily as training converges toward a solution that reconciles data fidelity with geotechnical consistency. Overall, L_{phy} drops from 0.0059 to 0.0005, corresponding to a 91.52% reduction. This marked reduction confirms that physics-constrained regularization effectively suppresses non-physical behaviors and yields deformation predictions that are both reliable and mechanistically consistent across

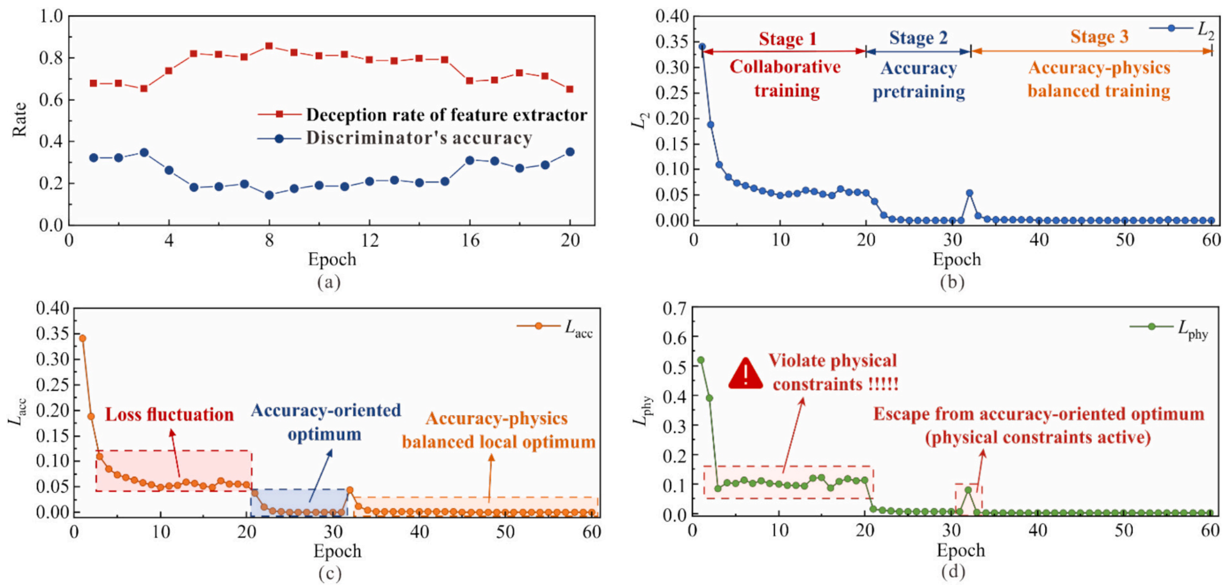

Fig. 10. Three-stage training process of the PL-DAST framework: (a) Variation of discriminator accuracy and feature extractor deception rate in Stage 1; (b) Evolution of total loss; (c) Evolution of accuracy loss; (d) Evolution of physics-constrained loss.

excavation stages.

In summary, the proposed three-stage training strategy offers distinct technical advantages. First, progressive optimization stabilizes physics-constrained learning and mitigates the convergence instability commonly observed in physics-constrained spatiotemporal deformation prediction model. Second, the staged design improves cross-project generalization by sequentially prioritizing domain alignment, accuracy-driven fitting, and physics-constrained refinement. Third, embedding soil loss constraints directly into training reduces physically implausible predictions and enhances interpretability beyond purely data-driven approaches. Collectively, these design features provide a robust and transferable learning paradigm for reliable spatiotemporal deformation prediction in deep foundation pits.

3.3. PL-DAST performance analysis

This subsection presents a comprehensive analysis of the experimental results to substantiate the advantages of the framework. The performance is evaluated through multiple dimensions, including the quality of learned domain-invariant features, cross-project generalization capability, and the verification of physical consistency based on soil loss theory.

3.3.1. Generation and evaluation of DIFs

Following training of the PL-DAST framework, the DIFs learned by the feature extractor require rigorous validation to substantiate the framework's cross-project transferability. Although Section 3.2.2 indicates that the feature extractor can effectively confuse the domain discriminator during adversarial learning, this qualitative evidence alone is insufficient. A quantitative evaluation is therefore needed to determine the extent to which the learned DIFs reduce cross-domain distributional discrepancies and support robust deformation prediction under unseen engineering conditions.

To evaluate cross-project feature alignment, this study combines qualitative visualization with quantitative discrepancy measurement. Specifically, t-SNE is used to visualize the distributional overlap of the original inputs and the corresponding DIFs, whereas the SWD provides a rigorous metric of cross-domain alignment. For a fair comparison, 100 samples are randomly selected from each of the three projects. Both the original 302-dimensional inputs and the extracted DIFs are then projected into a 3-dimensional space using t-SNE to enable direct inspection of clustering behavior. In parallel, SWD is computed by projecting the 201-dimensional representations onto 200 randomly sampled 1-dimensional subspaces and averaging the resulting Wasserstein distances. Together, these analyses quantify the extent to which PL-DAST reduces cross-project distributional discrepancies, thereby validating the encoder's ability to learn transferable representations.

Fig. 11 compares the distributions of the original inputs and the learned DIFs in a 3-dimensional embedding space. In Fig. 11(a), the original data form three project-specific clusters, indicating pronounced domain shifts that would hinder cross-project transfer. By contrast, the DIFs extracted after PL-DAST training in Fig. 11(b) exhibit substantially greater overlap and a more uniform spatial distribution. The weakened clustering and blurred project boundaries indicate effective feature alignment, confirming that the encoder suppresses project-dependent bias and yields transferable representations for downstream deformation prediction.

Table 5 quantifies the SWD between project pairs before and after feature extraction, complementing the t-SNE visual evidence. Results show that the mean SWD decreases from 0.127 for the original data to 0.112 after feature extraction, yielding a reduction rate of 11.81%. This indicates that the DANN feature extractor effectively reduces cross-project distributional discrepancies. Similar reductions are observed across all project pairs (10.91–12.05%), suggesting that the adversarial objective yields consistent, globally balanced alignment rather than pair-specific optimizations. Overall, these results confirm that the learned DIFs provide a transferable representation space, supporting reliable cross-project deformation prediction.

3.3.2. Cross-project generalization performance of PL-DAST

Building on the validation of DIF learning, this section evaluates the predictive performance of the PL-DAST framework under cross-project scenarios. Because excavation sites differ markedly in geological conditions and construction environments, such testing is critical for assessing real-world transferability. The results are presented below.

The PL-DAST framework exhibits outstanding cross-project generalization, delivering high prediction accuracy for both deformation prediction and physics-consistent inference in a previously unseen excavation project. Quantitative evaluation results are summarized in Table 6. PL-DAST achieves a remarkable R^2_{acc} of 0.996 on the training set and 0.993 on the test set, confirming effective transfer of learned deformation patterns to new projects. For the physical metric prediction,

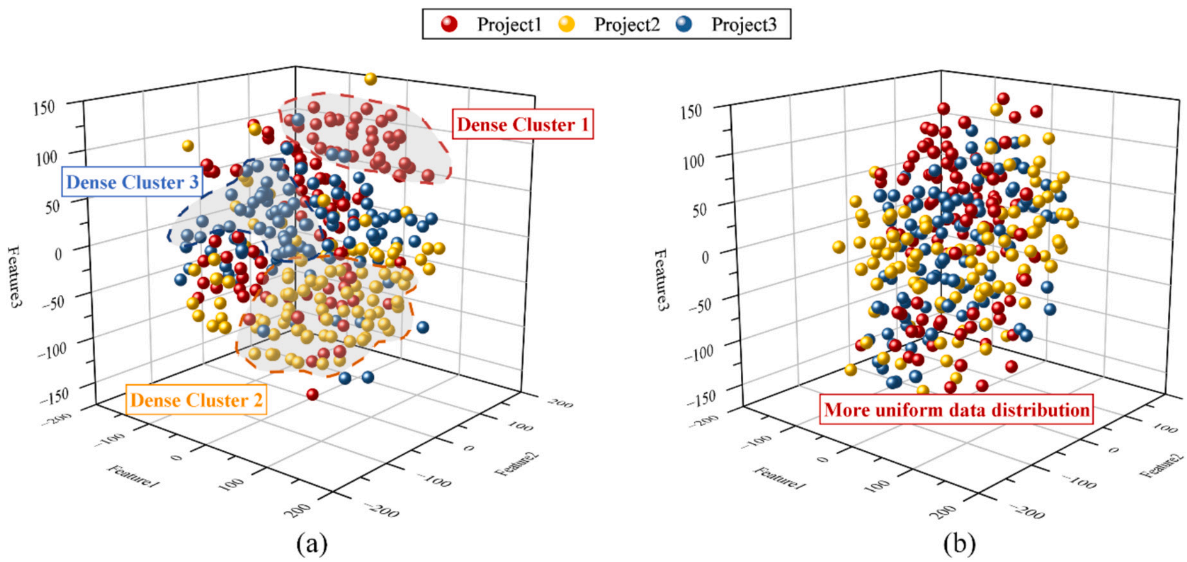

Fig. 11. Feature distribution alignment before and after domain-invariant feature learning: (a) Original data; (b) DIFs.

Table 5
Comparison of SWD between the original data and corresponding DIFs across project pairs.

Project pair	Original data	DIFs	Reduction rate (%)
Projects 1 & 2	0.166	0.146	12.05
Projects 1 & 3	0.159	0.140	11.95
Projects 2 & 3	0.055	0.049	10.91
Mean	0.127	0.112	11.81

Table 6
Cross-project prediction performance evaluation of PL-DAST on training and independent test set.

Dataset	R^2_{acc}	MSE_{acc}	R^2_{phy}	MSE_{phy}
Training set	0.996	1.274	0.990	0.0011
Test set	0.993	1.241	0.968	0.0009

the model obtains an R^2_{phy} of 0.990 (training set) and 0.968 (test set), showing slightly reduced performance in estimating the soil loss ratio m. This performance gap likely reflects cross-project differences in support configurations and soil-structure response, which increase the difficulty of transferring physical parameter inference. Nevertheless, the framework preserves $R^2_{phy} > 0.95$ on the unseen project, supporting its ability to generalize physical knowledge across heterogeneous excavation settings. These results collectively demonstrate that the PL-DAST delivers accurate and physically credible predictions, providing a reliable basis for deformation-informed risk assessment and excavation safety management.

3.3.3. Spatiotemporal deformation prediction performance of PL-DAST

The results in Section 3.3.2 confirm that PL-DAST achieves strong cross-project generalization. The next step is to assess its practical value for excavation monitoring by examining whether the framework can reproduce the spatiotemporal evolution of deformation and soil loss at the Lantian Road Station, which is treated as an unseen target site. Visual validation is conducted from two complementary perspectives: the spatial distribution of predicted responses and their temporal evolution.

(1) Spatial distribution.

(a) PL-DAST achieves spatially reliable predictions overall, but its errors remain spatially heterogeneous and concentrate in mechanically sensitive zones. Fig. 12 presents the spatial distribution of the TAE for retaining wall deformation and ground surface settlement across excavation sections. For retaining wall deformation, errors are strongly depth-dependent and form a distinct high error band between 16 and 28 m (Fig. 12(a)). This interval lacks steel strut support, so the wall must resist lateral earth pressure largely independently. The resulting complex soil-structure interaction increases deformation uncertainty and reduces prediction accuracy. For ground surface settlement, errors cluster primarily in the horizontal direction, with larger TAE values observed 5–30 m from the excavation edge (Fig. 12(b)). This region is subject to strong excavation-induced disturbance and exhibits settlement responses with pronounced spatial dependence. Overall, the spatial patterns of TAE are consistent with excavation geometry and support configuration, indicating that prediction difficulty is highest in regions governed by strong local deformation mechanisms.

(b) PL-DAST reproduces excavation-induced deformation with high spatial fidelity, demonstrating its ability to generalize deformation mechanisms in an unseen project.

To evaluate its predictive performance, three representative excavation stages are selected from Section 3 of the test set: an early stage (Day 5), an intermediate stage (Day 14), and a late stage (Day 22). As shown in Fig. 13, PL-DAST predictions for retaining wall deformation and ground surface settlement closely match field measurements, particularly in their spatial distribution patterns. Prediction errors increase with excavation depth, as soil-structure interactions become more complex. Nevertheless, the deviations remain within engineering-acceptable limits. This robustness is enabled by the complementary roles of the NAPL and DAGG modules. NAPL learns node-adaptive parameters to capture location-specific deformation behavior at each monitoring point. DAGG updates the spatial dependency graph to infer latent correlations among distant nodes. Together, these modules allow PL-DAST to model deformation evolution as a spatially coupled process rather than as independent pointwise time series.

(2) Temporal evolution.

PL-DAST captures the time-dependent deformation response well, although prediction uncertainty increases during excavation stages associated with intensified soil-structure interaction. As shown in Fig. 14 (a), the spatially averaged MAE (SA-MAE) for retaining wall deformation at Section 3 of Lantian Road Station increases steadily as excavation progresses. This continuous increase signifies a gradual decline in predictive accuracy as excavation deepens and deformation mechanisms become increasingly complex. In contrast, Fig. 14(b) presents a three-stage evolutionary pattern ("increase-decrease-stabilization") in the SA-MAE values of ground surface settlement. A pronounced peak

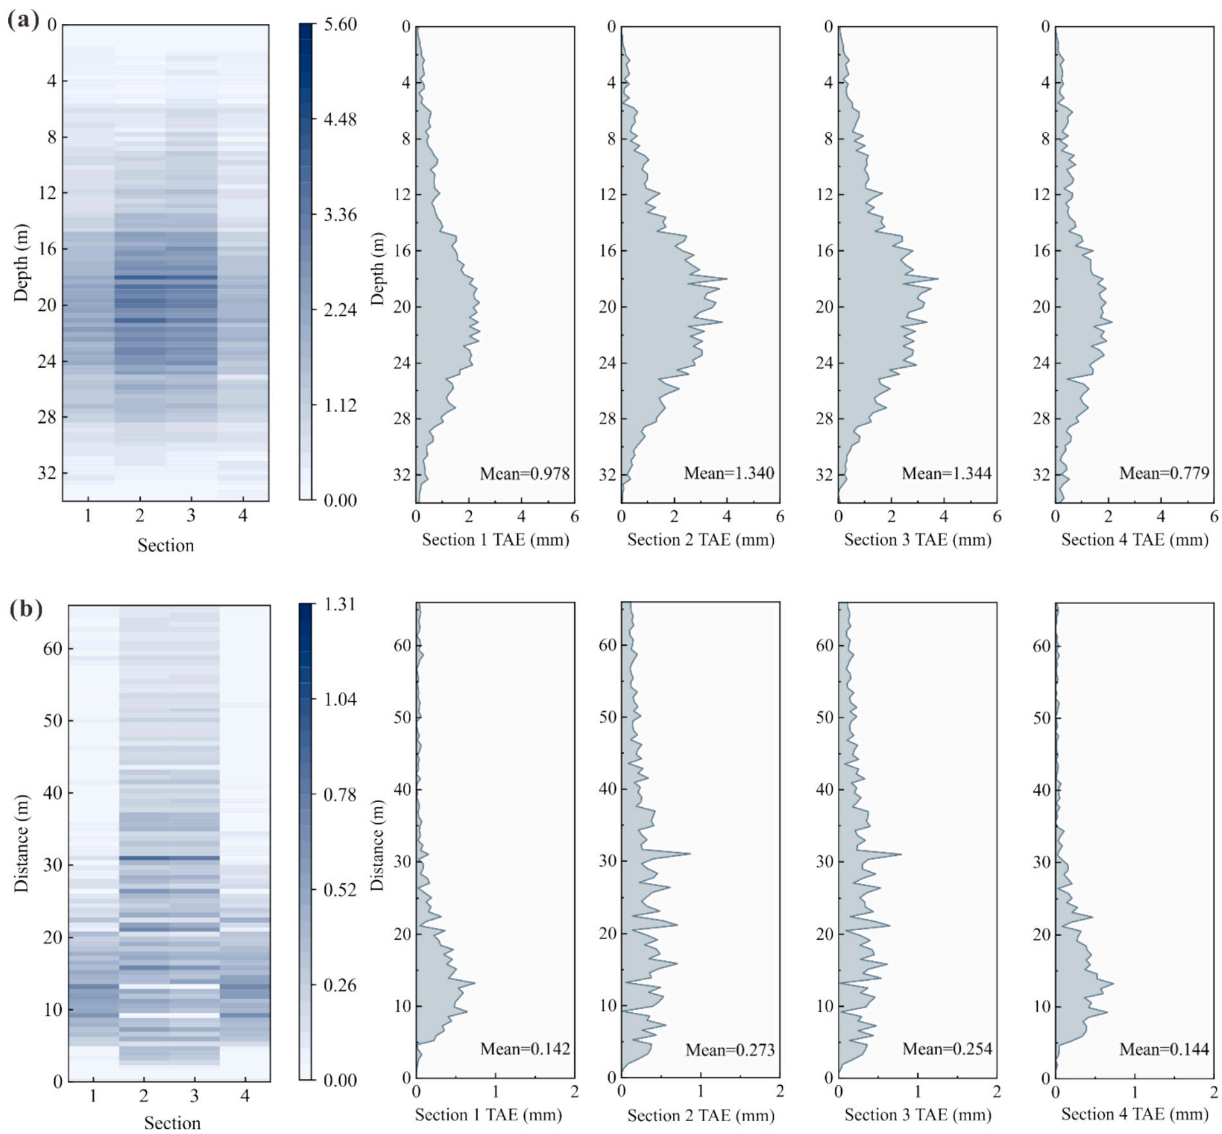

Fig. 12. Spatial distribution of TAE for PL-DAST across excavation sections: (a): Retaining wall deformation; (b) Ground surface settlement. (Note: The "distance" in Figs. 12 and 13 represent the distance to the excavation boundary.).

appears during Days 7–10, likely reflecting a period of accelerated settlement and intensified stress redistribution, which together increase prediction uncertainty.

Overall, these results highlight a pronounced temporal dependence in excavation-induced deformation, with different error evolution patterns across deformation modes. Retaining wall deformation shows near-linear accumulation of uncertainty, whereas ground surface settlement displays nonlinear and phase-dependent fluctuations. These dynamics underscore the importance of adaptive monitoring and responsive control strategies to manage time-varying risks throughout deep foundation pit excavation.

3.3.4. Physical consistency verification via soil loss theory

The PL-DAST framework exhibits strong physical consistency in predicting the soil loss ratio m at the Section 3 of Lantian Road Station, with predictions closely matching measured values throughout the excavation period. As shown in Fig. 15(a), PL-DAST reproduces the temporal evolution of m with high fidelity ($R^2 = 0.989$), demonstrating that the model preserves physical plausibility under changing excavation conditions.

Notably, the error evolution in Fig. 15(b) closely follows the excavation sequence. Following the first excavation stage (Days 5–8), soil stress redistribution increases modeling complexity and leads to larger errors. The second excavation (Day 9) triggers a rapid error escalation, with the largest deviation occurring on day 10. This peak is consistent with the cumulative disturbance of the soil-structure system. During the subsequent stabilization period (Days 11–17), errors decrease and remain comparatively low under the restraining effect of support installation. The third excavation (Day 18) then triggers renewed fluctuations, which persist until completion. These stage-dependent patterns indicate that the soil loss ratio m evolves dynamically with excavation activities, rather than remaining constant. This observation supports the core design of PL-DAST, which reformulates soil loss theory into a time-varying, physics-constrained constraint for cross-project deformation prediction.

4. Discussion

To quantify the respective contributions of LLM-derived textual embeddings, physics-constrained regularization, and domain-adversarial feature alignment in PL-DAST, a series of comparative experiments are conducted on the same dataset. First, an ablation study is

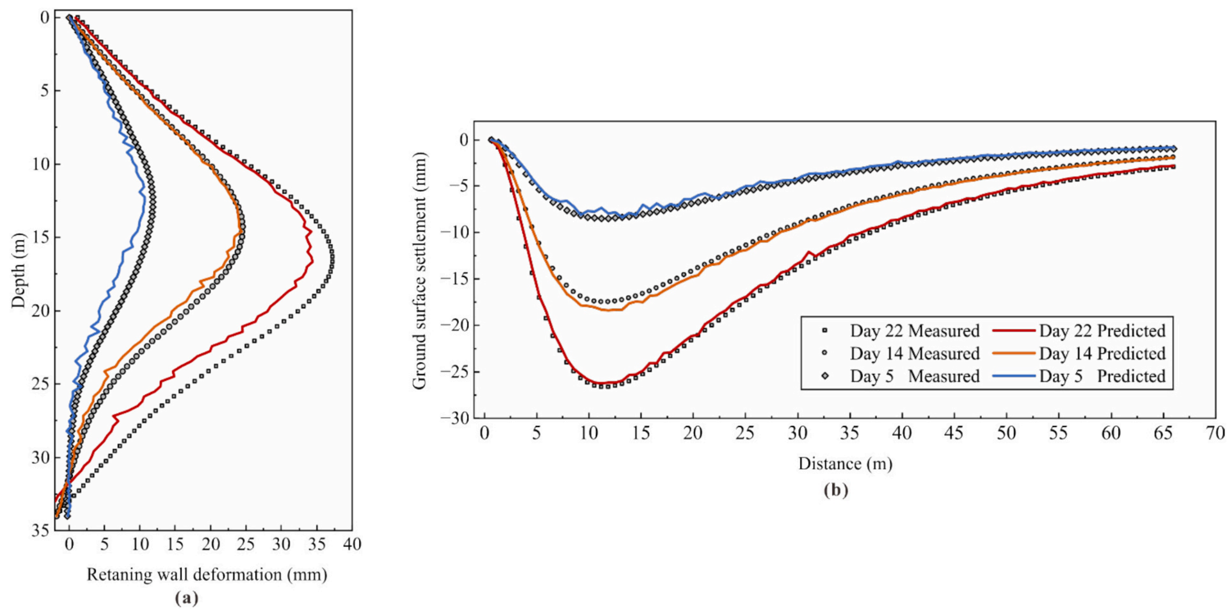

Fig. 13. Comparison of measured and PL-DAST-predicted excavation-induced deformations at Section 3 of the Lantian Road Station: (a): Retaining wall deformation; (b) Ground surface settlement.

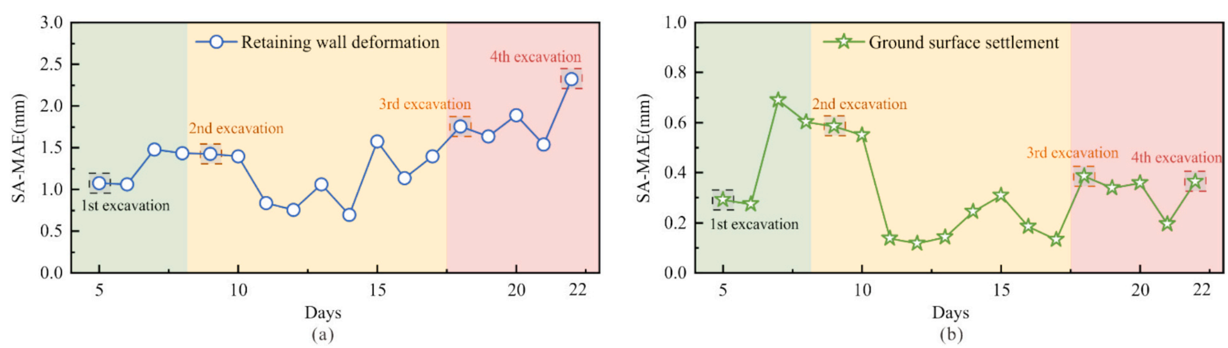

Fig. 14. Temporal evolution of SA-MAE for PL-DAST predictions at Section 3 of Lantian Road Station.

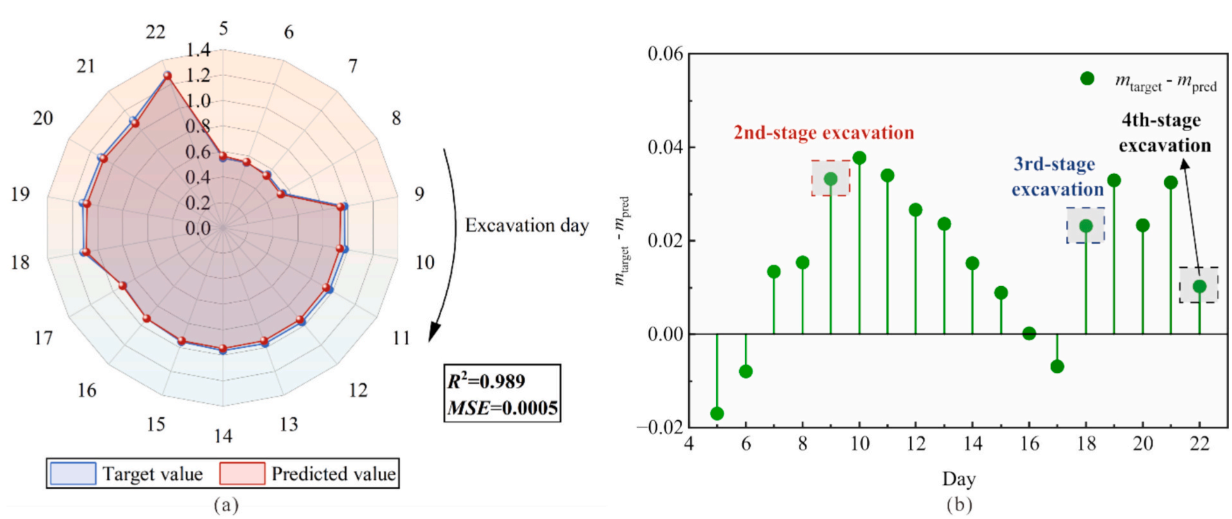

Fig. 15. Prediction results of soil loss ratio m by PL-DAST at Section 3 of Lantian Road Station: (a) Temporal evolution of predicted and measured values; (b) Temporal variation of prediction error.

performed by removing the LLM-based text inputs to isolate the effect of geological semantic features on representation learning and cross-project transfer. Second, the role of physical constraints is evaluated by varying the physics-based regularization and examining changes in predictive accuracy and physical consistency. Finally, PL-DAST is benchmarked against three widely used time-series forecasting algorithms to further highlight its predictive advantages. The results are summarized below.

(1) LLM-derived textual inputs significantly enhance both predictive accuracy and physical consistency in PL-DAST, particularly for retaining wall deformation. To isolate the contribution of text augmentation, this study compares models trained with and without LLM-derived semantic features under identical hyperparameter settings. The variant that does not use LLM-generated textual inputs is referred to as Physics-DAST. As demonstrated in Table 7, incorporating textual embeddings improves overall prediction accuracy: the prediction accuracy R^2_{acc} increases from 0.978 to 0.993, and the physical consistency R^2_{phy} rises from 0.958 to 0.968. Notably, the improvement is most pronounced for retaining wall deformation, where R^2 rises substantially from 0.931 to 0.977 (Fig. 16).

These improvements are attributable to the LLM's ability to encode geological investigation reports into semantic embeddings that supply the feature extractor with site-specific geological priors. By injecting geotechnical context into representation learning, PL-DAST learns deformation patterns that are both more transferable across projects and better aligned with physical behavior. Overall, the results show that integrating textual geological evidence enhances predictive accuracy and physical credibility, thereby improving the reliability and practical utility of deformation forecasting for deep foundation pit safety management.

(2) A precisely calibrated multimodal fusion weight strictly dictates the prediction performance of the PL-DAST framework. To systematically evaluate the robustness of the multimodal integration process, this study performs a targeted sensitivity analysis by varying the fusion weight coefficient while holding all other hyperparameters strictly constant. This coefficient controls the relative contribution of the LLM-derived geological embeddings during fusion with the numerical monitoring data. As demonstrated in Table 8, the predictive accuracy exhibits a highly nonlinear sensitivity to this parameter. Specifically, the optimal prediction performance is tightly concentrated within a narrow feasible region between 0.11 and 0.12, and outside this specific interval, the predictive accuracy drops noticeably.

This pronounced sensitivity arises from the inherent informational heterogeneity of the multimodal inputs. Although both inputs are projected into an equivalent dimensional space, the geological embeddings supply soft geotechnical knowledge, whereas the monitoring sequences dictate hard deterministic spatiotemporal dynamics. When the fusion weight diverges from the optimal range, this integration becomes severely unbalanced. An excessively large coefficient causes dense semantic priors to act as representational noise that overshadows precise numerical constraints and distorts actual deformation trajectories. Conversely, a coefficient that is too small prevents the PL-DAST framework from assimilating sufficient geological evidence for effective cross-project domain adaptation. Therefore, a precise mathematical alignment is strictly imperative to reconcile the soft geological knowledge with the hard numerical evidence, justifying the use of Bayesian optimization to ensure robust model transferability.

(3) Physics-constrained regularization substantially enhances the physical consistency of PL-DAST with only a marginal cost to predictive accuracy. To quantify the influence of physical regularization, the weight β is assigned to the physics loss L_{phy} from 0 to 0.9 in steps of 0.1, while holding all other hyperparameters constant. To account for stochastic variability in optimization, each setting is repeated 10 times and results are reported as averages.

As shown in Fig. 17, the accuracy metric R^2_{acc} exhibits only a slight decline as β increases, dropping from 0.997 at $\beta = 0$ to 0.981 at $\beta = 0.9$ (a 1.60% reduction). Specifically, the introduction of physical constraints exerts a differential impact on the two deformation modes. As β increases from 0 to 0.9, the R^2 for retaining wall deformation decreases noticeably from 0.995 to 0.910, whereas the R^2 for ground surface settlement declines more moderately from 0.995 to 0.969. This divergence suggests that the retaining wall deformation, which follows a relatively consistent structural pattern, is more sensitive to the imposition of the soil loss constraint. While the physics-constrained loss introduces a trade-off that slightly reduces the numerical fitting accuracy, particularly for the retaining wall, it is essential for enforcing the mechanistic coupling between the wall and the soil.

In contrast, the physics-consistency metric R^2_{phy} improves markedly, rising from 0.684 at $\beta = 0$ to a maximum of 0.947 at $\beta = 0.8$ (a 38.45% improvement). These results reveal a clear but manageable trade-off between numerical fidelity and physical plausibility. Importantly, introducing the physical constraint suppresses physically implausible solutions, reduces overfitting to project-specific signals, and guides optimization toward deformation trajectories that remain consistent with soil loss theory. This mechanism improves the robustness and engineering credibility of PL-DAST, strengthening its suitability for cross-project deformation prediction.

(4) PL-DAST consistently outperforms mainstream time-series forecasting baselines in cross-project deformation prediction, while achieving the strongest physical consistency. PL-DAST is benchmarked against the Gated Recurrent Unit (GRU), Long Short-Term Memory (LSTM) network, and Transformer under identical experimental settings. All models are trained using the same multimodal inputs (including monitoring data and geological text embeddings) and are equipped with the same physics-based constraint derived from soil loss theory. Hyperparameters are tuned via Bayesian optimization with 100 iterations to ensure fair comparison.

As summarized in Table 9, PL-DAST achieves the best overall predictive accuracy, with R^2_{acc} values of 0.996 (training) and 0.993 (test), respectively. More importantly, it delivers the highest level of physical plausibility, reaching R^2_{phy} values of 0.990 (training) and 0.968 (test). This advantage stems primarily from the AGCRN backbone, where the NAPL module learns node-specific parameters to capture spatial heterogeneity, and the DAGG module adaptively infers latent spatial dependencies across monitoring locations. Together, these mechanisms enable PL-DAST to exploit spatial-temporal coupling more effectively than purely sequential architectures, while integrating multimodal features and physics-constrained regularization in a unified framework.

The prediction error distributions in Fig. 18 further support this advantage: PL-DAST exhibits the narrowest and most sharply centered distributions around zero, indicating higher precision and stronger robustness than GRU, LSTM, and Transformer. Overall, PL-DAST achieves the intended balance between numerical accuracy and physical plausibility ($R^2_{phy} > 0.95$), supporting its practical value for cross-project deformation forecasting, early warning, and risk-informed safety management in deep foundation pit engineering.

Table 7
Comparison of prediction accuracy and physical consistency of PL-DAST with and without textual input on test set.

Model	R^2_{acc}	MSE_{acc}	R^2_{phy}	MSE_{phy}
PL-DAST (with textual input)	0.993	1.241	0.968	0.0009
Physics-DAST (without textual input)	0.978	3.638	0.958	0.0012

5. Conclusions and future work

To overcome the limited transferability and weak physical consistency of conventional DL approaches for deep foundation pit deformation prediction, this paper presented PL-DAST, a physics-constrained and LLM-driven domain-adversarial spatiotemporal framework for cross-project forecasting. Its key contribution is a unified encoder-decoder design that integrates geological semantics, domain invariant

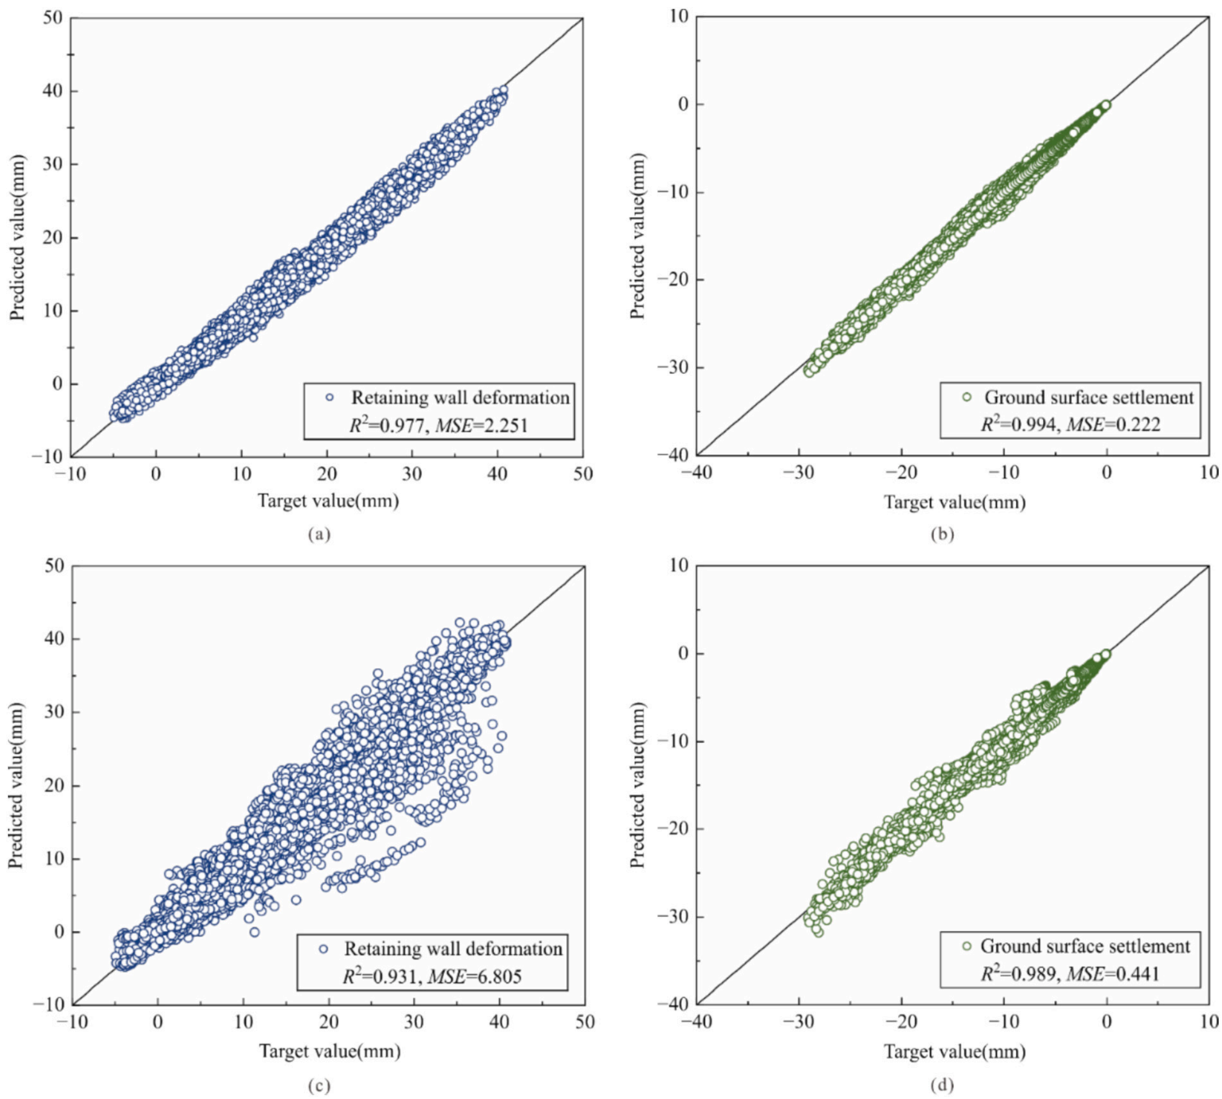

Fig. 16. Effect of LLM-derived textual inputs on deformation prediction: (a) and (c): Retaining wall deformation predicted by PL-DAST (with LLM-based text embeddings) and Physics-DAST (without textual input), respectively; (b) and (d): Ground surface settlement predicted by PL-DAST and Physics-DAST, respectively.

Table 8
Sensitivity analysis of PL-DAST prediction performance in test set to the fusion weight coefficient of the geological correction term.

Weight	R^2_{acc}	MSE_{acc}	R^2_{phy}	MSE_{phy}
0	0.988	2.026	0.866	0.004
0.1	0.725	46.406	0.932	0.002
0.11868	0.993	1.241	0.968	0.0009
0.2	0.685	53.272	0.812	0.005
0.3	0.627	63.075	0.845	0.004
0.4	0.575	71.695	0.833	0.005
0.5	0.730	45.545	0.829	0.005
0.6	0.759	40.622	0.780	0.006
0.7	0.763	40.006	0.804	0.006
0.8	0.714	48.276	0.771	0.007
0.9	0.682	53.740	0.833	0.005
1	0.653	58.688	0.791	0.006

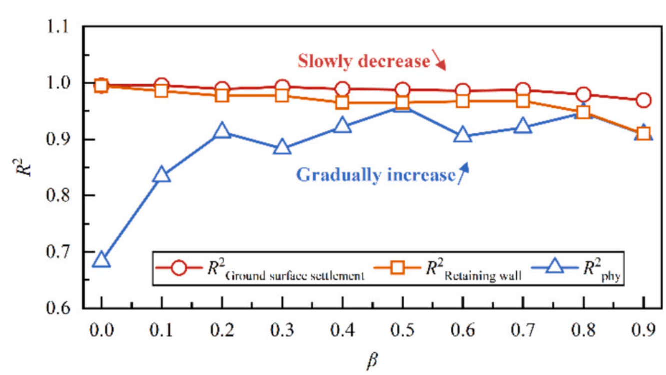

Fig. 17. Sensitivity of PL-DAST prediction performance to the weight β of physics-constrained loss.

representation learning, and geotechnical constraints in end-to-end training. The contributions are twofold: (a) From a theoretical perspective, PL-DAST introduces an LLM-enhanced domain-adversarial encoder that leverages Qwen3 to transform unstructured geological investigation reports into semantic embeddings and learns domain invariant features to align heterogeneous project data. The decoder couples AGCRN with NAPL and DAGG to model spatiotemporal dependencies, and incorporates soil loss theory as a physics-constrained regularizer. This design improves predictive accuracy while enforcing mechanistic consistency between retaining wall deformation and ground surface settlement. (b) From a practical standpoint, PL-DAST

Table 9
Comparison of predictive accuracy and physical consistency among four time series forecasting algorithms.

Model	Training set				Test set			
	R^2_{acc}	MSE_{acc}	R^2_{phy}	MSE_{phy}	R^2_{acc}	MSE_{acc}	R^2_{phy}	MSE_{phy}
PL-GRU	0.992	2.394	0.974	0.0026	0.980	3.370	0.921	0.0022
PL-LSTM	0.994	1.805	0.973	0.0027	0.981	3.157	0.921	0.0022
PL-Transformer	0.994	1.704	0.948	0.0053	0.986	2.398	0.932	0.0019
PL-DAST	0.996	1.274	0.990	0.0011	0.993	1.241	0.968	0.0009

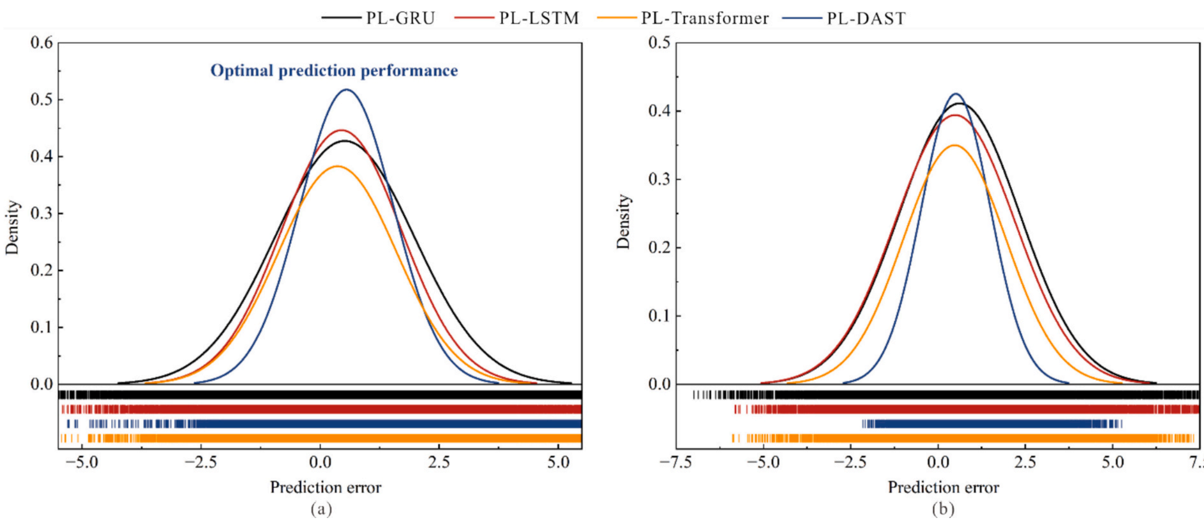

Fig. 18. Normal distribution curves of prediction errors for four time series forecasting algorithms on: (a) Training set; (b) Test set.

offers a practical and physics-consistent prediction system for deep excavation monitoring. By delivering reliable cross-project forecasts of key deformation responses, the framework can directly support deformation-based risk assessment, early warning, and risk-informed safety management in deep foundation pit practice.

Through rigorous validation on three real-world deep foundation pit projects, PL-DAST yields four main findings that conclusively demonstrate its superiority: (1) LLM-derived geological embeddings show strong cross-project similarity (minimum cosine similarity = 0.895) while preserving site-specific distinctions. Incorporating these text features substantially improves both predictive accuracy and physical plausibility relative to the text-free Physics-DAST baseline, underscoring the value of LLM-based encoding of geological reports. (2) Physics-constrained regularization markedly strengthens physical consistency: L_{phy} decreases by 91.52% (0.0059 to 0.0005), and R^2_{phy} increases from 0.684 to 0.947 (38.45%) with only a 1.60% reduction in R^2_{acc}. (3) The learned DIFs exhibit improved cross-project alignment, evidenced by greater overlap in feature space and an 11.81% reduction in average SWD (0.127 to 0.112). (4) PL-DAST achieves consistently high accuracy and physical consistency ($R^2_{acc} = 0.996/0.993$ and $R^2_{phy} = 0.990/0.968$ for training/test), supporting its use for intelligent safety monitoring and risk management in deep foundation pit engineering. Ultimately, this framework successfully bridges the gap between purely data-driven models and geotechnical mechanics, establishing a robust new paradigm for reliable cross-project deformation control in complex urban underground environments.

Although the proposed PL-DAST framework demonstrates strong accuracy and cross-project generalization, three limitations remain. First, the Qwen3 LLM is not fine-tuned on domain-specific foundation pit texts, which may restrict its ability to capture specialized geological terminology and reduce embedding quality. Second, the training data is drawn from two projects, limiting sample diversity and coverage of real-world conditions. Third, retaining wall deformation predictions occasionally show local jaggedness, indicating insufficient enforcement of spatial continuity. Future work should therefore fine-tune LLMs with domain-specific corpora, expand training to a broader range of projects and refine the physics-constrained loss to explicitly regularize spatial smoothness. These improvements are expected to further enhance reliability and practical utility in engineering applications.

CRediT authorship contribution statement

Yue Pan: Writing – review & editing, Writing – original draft, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition. **Zonghang He:** Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Jin-Jian Chen:** Writing – review & editing, Supervision, Investigation, Funding acquisition. **Xiang Li:** Writing – review & editing, Investigation, Funding acquisition. **Jianfeng Zhou:** Writing – review & editing, Validation, Investigation.

Declaration of competing 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.

Acknowledgement

This work was substantially supported by the National Key Research and Development Program of China (No. 2023YFC3009402), Shanghai Chenguang Program (No. 24CGA10), Young Talent Support Program of the Chinese Society for Rock Mechanics and Engineering (CSMRE-YESS20251109), and Shanghai Jiao Tong University AI for Engineering Initiative (No. 26×010100016).

Data availability

Data will be made available on request.

References

[1] X. Wang, Y. Pan, M. Li, J. Chen, A novel data-driven optimization framework for unsupervised and multivariate early-warning threshold modification in risk assessment of deep excavations, Expert Syst. Appl. 238 (2024) 121872, https://doi.org/10.1016/j.eswa.2023.121872.

[2] X. Li, Y. Pan, L. Zhang, J. Chen, Dynamic and explainable deep learning-based risk prediction on adjacent building induced by deep excavation, Tunn. Undergr. Space Technol. 140 (2023) 105243, https://doi.org/10.1016/j.tust.2023.105243.

[3] C. Yang, C. Wang, B. Wu, F. Zhao, J.-S. Fan, L. Zhou, Settlement estimation during foundation excavation using pattern analysis and explainable AI modeling, Autom. Constr. 166 (2024) 105651, https://doi.org/10.1016/j.autcon.2024.105651.

[4] C. Yang, C. Wang, F. Zhao, B. Wu, J.-S. Fan, Y. Zhang, Smart virtual sensing for deep excavations using real-time ensemble graph neural networks, Autom. Constr. 172 (2025) 106040, https://doi.org/10.1016/j.autcon.2025.106040.

[5] Y. Pan, J. Qin, L. Zhang, W. Pan, J.J. Chen, A probabilistic deep reinforcement learning approach for optimal monitoring of a building adjacent to deep excavation, Comput. Aided Civ. Inf. Eng. 39 (5) (2024) 656–678, https://doi.org/10.1111/mice.13021.

[6] X. Wang, Y. Pan, J. Chen, M. Li, A spatiotemporal feature fusion-based deep learning framework for synchronous prediction of excavation stability, Tunn. Undergr. Space Technol. 147 (2024) 105733, https://doi.org/10.1016/j.tust.2024.105733.

[7] H. Fan, L. Li, S. Zhou, M. Zhu, M. Wang, Prediction and risk assessment of lateral collapse in deep foundation pits using machine learning, Autom. Constr. 171 (2025) 106011, https://doi.org/10.1016/j.autcon.2025.106011.

[8] Y. Pan, W. He, J.-J. Chen, Spatiotemporal deep learning for multi-attribute prediction of excavation-induced risk, Autom. Constr. 171 (2025) 105964, https://doi.org/10.1016/j.autcon.2025.105964.

[9] J. Zhao, Z. Tan, R. Yu, Z. Li, X. Zhang, P. Zhu, Deformation responses of the foundation pit construction of the urban metro station: a case study in Xiamen, Tunn. Undergr. Space Technol. 128 (2022) 104662, https://doi.org/10.1016/j.tust.2022.104662.

[10] H. Kuang, Z. Ai, G. Gu, H. Huang, Z. Qian, D. Zhou, Y. Zhai, H. Guo, A dynamic prediction model for maximum lateral displacement of foundation pit enclosure structure using multi-source fusion method, Comput. Geotech. 187 (2025) 107468, https://doi.org/10.1016/j.compgeo.2025.107468.

[11] C. Zhao, Y. Feng, W. Wang, Z. Niu, Mechanical properties and numerical analysis of underground continuous wall in underground grain silo foundation pit, Buildings 13 (2) (2023) 293, https://doi.org/10.3390/buildings13020293.

[12] M.-G. Li, Q.-Z. Xiao, N.-W. Liu, J.-J. Chen, Predicting wall deflections for deep excavations with servo struts in soft clay, J. Geotech. Geoenviron. Eng. 150 (1) (2024) 04023124, https://doi.org/10.1061/JGGEFK.GTENG-11347.

[13] J. Zhang, R. Xie, H. Zhang, Mechanical response analysis of the buried pipeline due to adjacent foundation pit excavation, Tunn. Undergr. Space Technol. 78 (2018) 135–145, https://doi.org/10.1016/j.tust.2018.04.026.

[14] L. Mu, M. Huang, Small strain based method for predicting three-dimensional soil displacements induced by braced excavation, Tunn. Undergr. Space Technol. 52 (2016) 12–22, https://doi.org/10.1016/j.tust.2015.11.001.

[15] H. Huang, Z. Liu, Y. Wang, H. Wang, Parametric modeling and safety simulation of pit excavation affecting adjacent tunnels based on BIM-FEM framework, Tunn. Undergr. Space Technol. 164 (2025) 106800, https://doi.org/10.1016/j.tust.2025.106800.

[16] Y. Pan, J. Qin, Y. Hou, J.-J. Chen, Two-stage support vector machine-enabled deep excavation settlement prediction considering class imbalance and multi-source uncertainties, Reliab. Eng. Syst. Saf. 241 (2024) 109578, https://doi.org/10.1016/j.ress.2023.109578.

[17] W.-S. Zhang, Y. Yuan, M. Long, R.-H. Yao, L. Jia, M. Liu, Prediction of surface settlement around subway foundation pits based on spatiotemporal characteristics and deep learning models, Comput. Geotech. 168 (2024) 106149, https://doi.org/10.1016/j.compgeo.2024.106149.

[18] S. Wang, B. Han, J. Jiang, N. Telyatnikova, Machine learning and FEM-driven analysis and optimization of deep foundation pits in coastal area: a case study in Fuzhou soft ground, Underground Space 22 (2025) 55–76, https://doi.org/10.1016/j.undsp.2024.11.001.

[19] X. Wang, Y. Pan, J.J. Chen, Digital twin with uncertainty-informed deep learning for prompt quantitative risk assessment of deep excavation, Comput. Aided Civ. Inf. Eng. (2025), https://doi.org/10.1111/mice.70049.

[20] X. Zhou, Y. Pan, J. Qin, J.-J. Chen, P. Gardoni, Spatio-temporal prediction of deep excavation-induced ground settlement: a hybrid graphical network approach considering causality, Tunn. Undergr. Space Technol. 146 (2024) 105605, https://doi.org/10.1016/j.tust.2024.105605.

[21] A. Dahal, L. Lombardo, Towards physics-informed neural networks for landslide prediction, Eng. Geol. 344 (2025) 107852, https://doi.org/10.1016/j.enggeo.2024.107852.

[22] Z. Zhang, Q. Pan, Z. Yang, X. Yang, Physics-informed deep learning method for predicting tunnelling-induced ground deformations, Acta Geotech. 18 (9) (2023) 4957–4972, https://doi.org/10.1007/s11440-023-01874-9.

[23] W. Ouyang, G. Li, L. Chen, S.-W. Liu, Machine learning-based soil–structure interaction analysis of laterally loaded piles through physics-informed neural networks, Acta Geotech. 19 (7) (2024) 4765–4790, https://doi.org/10.1007/s11440-023-02179-7.

[24] Y. Shi, M. Beer, Physics-informed neural network classification framework for reliability analysis, Expert Syst. Appl. 258 (2024) 125207, https://doi.org/10.1016/j.eswa.2024.125207.

[25] K. Luo, J. Zhao, Y. Wang, J. Li, J. Wen, J. Liang, H. Soekmadji, S. Liao, Physics-informed neural networks for PDE problems: a comprehensive review, Artif. Intell. Rev. 58 (10) (2025) 1–43, https://doi.org/10.1007/s10462-025-11322-7.

[26] J.-Z. Peng, Y. Hua, N. Aubry, Z.-H. Chen, M. Mei, W.-T. Wu, Data and physics-driven modeling for fluid flow with a physics-informed graph convolutional neural network, Ocean Eng. 301 (2024) 117551, https://doi.org/10.1016/j.oceaneng.2024.117551.

[27] J. Li, S. Zhang, X. Wang, Physics-informed neural network with fuzzy partial differential equation for pavement performance prediction, Autom. Constr. 171 (2025) 105983, https://doi.org/10.1016/j.autcon.2025.105983.

[28] H. Yan, H.-Y. Gou, F. Hu, Y.-Q. Ni, Y.-W. Wang, D.-C. Wu, Y. Bao, Smart control of bridge support forces using adaptive bearings based on physics-informed neural network (PINN), Autom. Constr. 168 (2024) 105790, https://doi.org/10.1016/j.autcon.2024.105790.

[29] A.M. Roy, R. Bose, V. Sundararaghavan, R. Arróyave, Deep learning-accelerated computational framework based on physics informed neural network for the solution of linear elasticity, Neural Netw. 162 (2023) 472–489, https://doi.org/10.1016/j.neunet.2023.03.014.

[30] R. Bose, A.M. Roy, Invariance embedded physics-infused deep neural network-based sub-grid scale models for turbulent flows, Eng. Appl. Artif. Intell. 128 (2024) 107483, https://doi.org/10.1016/j.engappai.2023.107483.

[31] A.M. Roy, S. Guha, A data-driven physics-constrained deep learning computational framework for solving von Mises plasticity, Eng. Appl. Artif. Intell. 122 (2023) 106049, https://doi.org/10.1016/j.engappai.2023.106049.

[32] A.M. Roy, S. Guha, V. Sundararaghavan, R. Arróyave, Physics-infused deep neural network for solution of non-associative Drucker–Prager elastoplastic constitutive model, J. Mech. Phys. Solids 185 (2024) 105570, https://doi.org/10.1016/j.jmps.2024.105570.

[33] J. Dagdelen, A. Dunn, S. Lee, N. Walker, A.S. Rosen, G. Ceder, K.A. Persson, A. Jain, Structured information extraction from scientific text with large language models, Nat. Commun. 15 (1) (2024) 1418, https://doi.org/10.1038/s41467-024-45563-x.

[34] P. Guo, H. Xue, J. Ma, J.C.P. Cheng, Advancing BIM information retrieval with an LLM-based query-domain-specific language and library code function alignment system, Autom. Constr. 178 (2025) 106374, https://doi.org/10.1016/j.autcon.2025.106374.

[35] H. Wu, J. Li, Q. Zhang, J. Tao, Z. Meng, Intelligent fault diagnosis of rolling bearings under varying operating conditions based on domain-adversarial neural network and attention mechanism, ISA Trans. 130 (2022) 477–489, https://doi.org/10.1016/j.isatra.2022.04.026.

[36] G. Li, Z. Wang, J. Gao, C. Xu, Y. Guo, D. Sun, X. Fang, Performance assessment of cross office building energy prediction in the same region using the domain adversarial transfer learning strategy, Appl. Therm. Eng. 241 (2024) 122357, https://doi.org/10.1016/j.applthermaleng.2024.122357.

[37] D. Liu, W. Zhang, K. Duan, J. Zuo, M. Li, X. Zhang, X. Huang, X. Liang, Intelligent prediction and optimization of ground settlement induced by shield tunneling construction, Tunn. Undergr. Space Technol. 160 (2025) 106486, https://doi.org/10.1016/j.tust.2025.106486.

[38] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. March, V. Lempitsky, Domain-adversarial training of neural networks, J. Mach. Learn. Res. 17 (59) (2016) 1–35, https://doi.org/10.1007/978-3-319-58347-1_10.

[39] T. Chen, X. Zhang, C. Wang, X. Yu, S. Wang, X. Chen, Domain adversarial neural network-based nonlinear system identification for helicopter transmission system, Nonlinear Dynamics 111 (16) (2023) 14695–14711, https://doi.org/10.1007/s11071-023-08657-7.

[40] B. Dai, G. Frusque, T. Li, Q. Li, O. Fink, Smart filter aided domain adversarial neural network for fault diagnosis in noisy industrial scenarios, Eng. Appl. Artif. Intell. 126 (2023) 107202, https://doi.org/10.1016/j.engappai.2023.107202.

[41] D. Hu, V.J. Gan, R. Zhai, Automated BIM-to-scan point cloud semantic segmentation using a domain adaptation network with hybrid attention and whitening (DawNet), Autom. Constr. 164 (2024) 105473, https://doi.org/10.1016/j.autcon.2024.105473.

[42] S. Ye, X. Song, M.J. Blunt, B. Bijeljic, Domain adaptation neural network based prediction of real-time drill bit tooth wear, Geoenergy Sci. Eng. 240 (2024) 213041, https://doi.org/10.1016/j.geoen.2024.213041.

[43] A. Yang, A. Li, B. Yang, B. Zhang, B. hui, B. Zheng, B. Yu, C. Gao, C. Huang, C. Lv, Qwen3 technical report, arXiv Preprint (2025), https://doi.org/10.48550/arXiv.2505.09388 arXiv:2505.09388.

[44] J. Liu, X. Hou, Excavation Engineering Handbook, China Architecture and Building Press, Beijing, 1997.

[45] L. Bai, L. Yao, C. Li, X. Wang, C. Wang, Adaptive graph convolutional recurrent network for traffic forecasting, Adv. Neural Inf. Proces. Syst. 33 (2020) 17804–17815, https://doi.org/10.5555/3495724.3497218.