Proc. of the European Safety and Reliability Conference (ESREL 2026)
Edited by José C. Matos, Paulo B. Lourenço, Michael Beer, Daniel Oliveira,
Edoardo Patelli, Hélder S. Sousa, Rui A. Silva and Monica Santamaria

doi: 10.3850/ESREL2026061419_esrel26-p29042-cd

From Multi-Source Monitoring to Trustworthy Decisions: Hybrid AI–Physics State-Space Models for Transportation Infrastructure

Ylli Murati

University of Minho, ISISE, ARISE, Department of Civil Engineering, Portugal, yllimuratti@gmail.com

Joaquim Tinoco

University of Minho, ISISE, ARISE, Department of Civil Engineering, Portugal, *jtinoco@civil.uminho.pt*

Luis Matos

ALGORITMI/LASI Center – Department of Information Systems, School of Engineering, University of Minho, Campus de Azurém, Guimarães, 4800-058, Braga, Portugal, EPMQ, CCG/ZGDV Institute, Campus de Azurém, Guimarães, 4804-533, Braga, Portugal, luis.matos@dsi.uminho.pt

Luis Magalhães,

ALGORITMI/LASI Research Center, Dep. Information Systems, Minho University, Guimarães, Portugal, lmagalhaes@dsi.uminho.pt

Jose Matos

University of Minho, ISISE, ARISE, Department of Civil Engineering, Portugal, jmatos@civil.uminho.pt

Steffan Davies

Spotlite-AETHRA, Rua Pedro Nunes, IPN, Bloco-C 3030-199, Coimbra, Portugal, steffan@spotlitedata.com

Transportation infrastructure monitoring is increasingly characterized by heterogeneous, multi-rate data streams spanning in-situ IoT sensing (e.g., accelerations, strains, inclinations, temperatures), geodetic measurements, and remote sensing (e.g., InSAR displacement time series). Converting such observations into uncertainty-aware, decision-support-oriented information remains challenging due to missingness, asynchronous sampling, non-stationary environmental effects, and the need to propagate uncertainty from measurement to maintenance action. This paper proposes a hybrid AI–physics state-space framework that unifies representation learning, probabilistic data assimilation, and unsupervised regime discovery for network-scale asset management.

A masked, self-supervised encoder first converts heterogeneous sensor and remote-sensing observations into a compact latent representation that can handle missing data and imperfect alignment between sources. These latent measurements are then fused using an asynchronous, multi-rate Kalman Filter (KF) defined on a physically interpretable latent state representation, where dynamic features such as displacement trends and velocity-like behaviour are implicitly encoded by the neural encoder, so that InSAR and in-situ sensors can be combined in a dynamically consistent way. To account for changing behavior over time and hidden deterioration patterns, the KF state trajectories are grouped using Bayesian non-parametric mixture models, which can discover the appropriate number of regimes automatically and assign each time period a probability of belonging to each regime. Overall, the framework outputs uncertainty-aware state estimates, forecasts, and regime probabilities, enabling probabilistic alarms and decision rules (e.g., exceedance probabilities and confidence-weighted warnings) instead of fixed threshold triggers. Collectively, the framework advances the state of practice from “monitoring and visualization” toward uncertainty-calibrated, explainable decision support for transportation infrastructure.

Keywords: Hybrid AI–physics modeling, state-space models, InSAR displacement time series, Kalman Filter (KF), Bayesian non-parametric clustering, uncertainty quantification.

1. Introduction

Geotechnical assets such as natural slopes, embankments, and cuttings are subject to progressive deformation driven by a combination of intrinsic material behaviour, hydrological loading, and environmental cycles. Satellite-based Interferometric Synthetic Aperture Radar (InSAR) now provides millimetric-precision displacement time series over large spatial domains, enabling continuous observation of surface deformation at network scale [1, 2]. However, converting such observations into actionable condition intelligence remains challenging due to the heterogeneity of data sources, the absence of labelled failure records, non-stationary environmental effects, and the need to propagate uncertainty from measurement to engineering decision.

Current practice commonly treats InSAR signals in isolation or applies rule-based threshold triggers that do not account for the interaction between environmental forcing and structural response. Approaches that integrate rainfall as an explicit driver of slope deformation, or that couple multi-directional InSAR responses within a unified dynamic model, remain limited in operational deployment [3, 4].

This paper addresses this gap by proposing a unified multi-modal AI–physics pipeline that jointly encodes InSAR horizontal displacement and precipitation data, filters the resulting latent representations through a Kalman-based state-space model, and classifies monitoring points into coherent dynamic response regimes using unsupervised Bayesian clustering. The framework is developed within the WISE (Wireless Transportation Infrastructure Safety Evaluation) project and is designed for extensibility: vibration sensors, temperature, strain gauges, geometric metadata, and historical failure records can be incorporated without redesigning the core architecture. Although the present study demonstrates the framework on a geotechnical slope-monitoring case study, the underlying state-space and probabilistic fusion architecture is designed to be asset-agnostic and extensible to other transportation infrastructure systems, including bridges, retaining walls, embankments, and structural health monitoring applications, as additional multi-modal datasets become available.

2. State of the Art

2.1 Remote Sensing and InSAR for Slope Monitoring

Over the past two decades, satellite InSAR and multi-temporal InSAR (MTInSAR) have become standard tools for detecting and characterising slope instabilities at regional and local scales [1, 5]. Casagli et al. [5] synthesised advances across three functional domains: landslide inventory mapping, deformation monitoring and early warning, and susceptibility and hazard assessment, establishing InSAR as central to all three. Scaioni et al. [6] formalised the multi-sensor perspective, mapping optical, LiDAR, SAR, and UAV technologies onto these functions and identifying multi-sensor integration as a critical research direction.

Operational monitoring frameworks combining satellite InSAR with ground-based sensing have been implemented for natural mountain slopes [7, 8] and transportation infrastructure [9]. Vaccari et al. [9] integrated satellite InSAR into ArcGIS-based decision support tools for network-level screening of slope instabilities, bridge settlements, and pavement distress. Despite these advances, existing frameworks remain domain-specific and rarely couple deformation monitoring with explicit environmental forcing within a unified probabilistic model.

2.2 Machine Learning for Slope Stability Assessment

Machine learning methods have been widely applied to slope stability classification, predominantly in supervised settings. Zhang and Wei [10] compared Gradient Boosted Models, SVM, XGBoost, and Random Forest on 188 slope cases, identifying GBM as superior for correlation with finite element model results. Bai et al. [11] evaluated eight ML models and identified ANN as the most accurate across infrastructure datasets. Hybrid approaches combining Random Forest with metaheuristic optimisation (PSO, firefly algorithms) and least-squares SVM have demonstrated strong predictive performance [12].

Tinoco et al. [13] developed ANN-based models using extensive UK railway inspection data, demonstrating robust stability classification when addressing class imbalance through oversampling; this was extended in [14] through comparison with RF and SVM. Tao [15] integrated FEM with SVM to account for rainfall and water-level variations in riverbank slopes, an early example of physics-informed hybrid modelling for slope assessment.

Unsupervised learning approaches, including K-means clustering, Gaussian Mixture Models, and anomaly detection, have been applied where labelled failure data are unavailable [16]. However, the integration of temporal dynamics through state-space modelling with unsupervised probabilistic clustering has not been widely demonstrated for slope monitoring applications.

2.3 State-Space Models and Bayesian Approaches

State-space models (SSMs) provide a principled framework for representing dynamic system behaviour under uncertainty, enabling recursive probabilistic state estimation as new observations arrive [17]. Kalman-based approaches have been applied to bridge monitoring and infrastructure deterioration modelling, with Bayesian extensions enabling uncertainty propagation and continuous model updating [18]. Von Rueden et al. [18] demonstrated that hybrid approaches combining physics-based simulation with data-driven learning improve robustness under heterogeneous data conditions. Miao et al. [19] applied a hybrid attention-LSTM and Hidden Markov Model to bridge condition assessment under sparse inspection data, confirming the utility of SSM-based formulations for infrastructure monitoring. The application of Bayesian non-parametric SSMs to geotechnical slope monitoring with multi-modal sensor fusion remains an open research direction that this paper addresses.

3. Proposed Framework

The proposed framework follows a four-stage pipeline: (i) multi-modal data loading and preprocessing, (ii) self-supervised multi-modal encoding, (iii) Kalman filter-based state estimation, and (iv) unsupervised Bayesian clustering and probabilistic inference. Figure 1 provides a schematic overview.

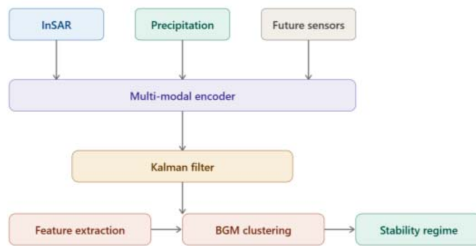

Fig. 1. Multi-modal AI–physics pipeline for slope stability classification.

3.1 Data Loading and Preprocessing

Raw sensor data are loaded from CSV files and aligned to a common temporal axis using the InSAR horizontal channel as the reference modality. Input normalisation is applied per modality as $(x - \mu) / (\sigma + 10^{-4})$ prior to encoding. A sliding window of length $W = 5$ timesteps with a step of 1 is applied to all dynamic modalities, producing $N = T - W + 1$ overlapping windows per monitoring point. Static modalities (geometry, metadata, historical failures) bypass windowing and are tiled across temporal positions after encoding. A base sliding window of length $W=5$ is used; when fewer timesteps are available, the encoder operates adaptively on the available sequence length (e.g., 2–4 timesteps).

The training dataset consists of $P = 671$ monitoring points with $T = 85$ timesteps.

3.2 Multi-Modal Encoder

The encoder maps each dynamic modality into a shared latent space of dimension $d = 32$. Three encoder architectures are deployed according to modality type:

CNN Encoder — used for InSAR displacement channels. A 1D convolutional layer (kernel size 5) processes each window of W timesteps. Features are averaged over the window dimension. A variance head outputs per-window uncertainty estimates $\sigma^2 \in \mathbb{R}^{N \times d}$.

LSTM Encoder — used for precipitation and temperature. Each window is processed sequentially; the hidden state at the final timestep serves as the window representation. A linear variance head produces uncertainty estimates.

MLP Encoder — used for static modalities. Projects a fixed-size feature vector to $\mathbb{R}^d$ and tiles it across all N temporal positions.

When both horizontal and vertical InSAR channels are available, an attention mechanism fuses them adaptively, (see Eq. (1)).

$$z_{InSAR} = \alpha \cdot z_h + (1 - \alpha) \cdot z_v, \tag{1}$$

where $\alpha = \sigma(W \cdot [z_h \parallel z_v])$. Precipitation modulates the InSAR latent representation through Feature-wise Linear Modulation (FiLM), (see Eq. (2)).

$$\tilde{z}_{InSAR} = z_{InSAR}(1+\gamma / (\sigma^2_{rain} + \varepsilon))+(\beta / (\sigma^2_{rain}+\varepsilon)) \tag{2}$$

where γ and β are learned linear transformations and uncertainty scaling suppresses unreliable precipitation signals. All modality representations are fused using inverse-variance weighting: $w_m = 1 / (\sigma^2_m + \varepsilon)$, giving higher weight to modalities with lower uncertainty.

The encoder is trained in a self-supervised manner by minimising the reconstruction loss, (see Eq. (3)).

$$L=MSE(\hat{x}_{InSAR}, x_{InSAR})+MSE(\hat{x}_{Rain}, x_{Rain}) \tag{3}$$

Mini-batch training (batch size 256) is used to manage memory requirements. Training used 100 epochs with Adam optimiser (lr = 10^{-3}), reducing reconstruction loss from ≈1.5 to ≈0.01.

3.3 Kalman Filter

A Kalman filter with identity state transition (F = I) and observation matrix (H = I) is applied to the fused latent sequence $z_{fused} \in \mathbb{R}^{P \times N \times d}$. The observation noise covariance $R_t = \mathrm{diag}(\sigma^2\_t)$ is set dynamically from the encoder's per-timestep uncertainty output, allowing the filter to discount unreliable encoder outputs automatically. The filter equations are ((see Eq. (4)(5)(6)).

$$\text{Predict: } \hat{x}_{\{t|t-1\}} = x_{\{t-1\}}, \quad P_{\{t|t-1\}} = P_{\{t-1\}} + Q \tag{4}$$

$$\text{Update: } S_t = P_{\{t|t-1\}} + R_t, \quad K_t = P_{\{t|t-1\}} \quad S_t^{-1} \tag{5}$$

$$x_t = \hat{x}_{\{t|t-1\}} + K_t\,(z_t - \hat{x}_{\{t|t-1\}}), \quad P_t = (I - K_t)\, P_{\{t|t-1\}} \tag{6}$$

The filter is fully vectorized over the P dimension, processing all monitoring points in parallel. Runtime on CPU is approximately 3 seconds for P = 671, N = 85, d = 32.

Since the state variables are learned latent features (including velocity- and trend-related information extracted by the encoder), an identity transition model (F = I) is adopted, allowing flexible temporal evolution without imposing rigid kinematic constraints.

3.4 Feature Extraction and Clustering

After Kalman filtering, a feature vector is computed for each monitoring point from the smoothed state sequence and encoder outputs. The feature set includes: statistical features (mean, variance, min, max, amplitude range, energy), kinematic features (trend, mean/max/min velocity, mean acceleration), and cross-modal features (Pearson correlation between InSAR and precipitation latent embeddings at lags τ = 0, 1, 2, 3 timesteps). Features are standardised and compressed to k principal components via PCA, with the final model using k = 3 based on clustering performance.

Unsupervised clustering uses a Bayesian Gaussian Mixture Model (BGM) with a Dirichlet process prior, which automatically determines the effective number of clusters up to a maximum of five. For inference on new monitoring points, the trained BGM computes, (see Eq. (7)).

$$p(c= k \mid f) = (\pi_k\, N(f; \mu_k, \Sigma_k)) / (\Sigma_j\, \pi_j\, N(f; \mu_j, \Sigma_j)) \tag{7}$$

where f is the PCA-compressed feature vector, π_k are mixture weights, and (μ_k, Σ_k) are cluster parameters. This provides soft probability assignments suitable for risk-based decision support.

4. Results and Discussion

4.1 Clustering Performance

Table 1 summarizes clustering results for different PCA configurations on the training set of 671 monitoring points.

Table 1. Clustering performance across configurations on P = 671 training points.

Method	PCA Components	Silhouette Score	Active Clusters
BGM	3	0.347	5
BGM	6	0.271	5
BGM	32	-0.097	5
HDBSCAN	2	0.318	4 + noise points

The best configuration, BGM with PCA compressed to 3 components, achieved a silhouette score of 0.347 with 5 active clusters. Visual inspection of the 2D cluster space reveals a characteristic arc structure in the InSAR data, reflecting the physical distribution of surface displacement across the monitored points: a continuum from stable zones (near-zero displacement) to actively deforming zones (large negative displacement), with a small population of anomalous points at extreme displacement values.

Table 2. Clustering evaluation metrics for the trained BGM model

Metric	Value
Silhouette Score	0.347
Davies-Bouldin Index	0.869
Mean Classification Entropy	0.236
Mean Confidence	90.57%

The arc structure is physically meaningful and expected in InSAR-based slope monitoring: it encodes the gradient from stable to unstable behaviour as a smooth manifold in latent space rather than as discrete, well-separated Gaussian clusters. The silhouette score reflects this geometry, as BGM Gaussian components approximate an arc rather than spherical clusters, and is consistent with results reported for similar unsupervised monitoring frameworks applied to geotechnical data [16].

From table 2 the Silhouette Score measures how similar each point is to its own cluster compared to neighbouring clusters. A score of 0.347 falls in the moderate range, which is expected for geotechnical displacement data that forms a physical continuum rather than discrete, well-separated groups. Values above 0.5 are typical for clearly separable synthetic datasets; for real-world continuous phenomena this range is consistent with meaningful cluster structure. Davies-Bouldin Index evaluates the average similarity between each cluster and its most similar neighbour. A value of 0.869 indicates moderate overlap between some cluster pairs. Mean Classification Entropy quantifies the average uncertainty across all assignments. A value of 0.236, close to zero, indicates that most points receive clearly dominant probability assignments, with limited ambiguity across the five regimes. Mean Confidence directly expresses the average maximum posterior probability across all 671 training points. A mean confidence of 90.57% with only moderate silhouette suggests that while some cluster boundaries are geometrically close, the Bayesian model assigns points decisively. Despite the geometric overlap of the 2D cluster space, the high confidence of regime assignment is a consequence of probabilistic assignment.

4.2 Inference on New Monitoring Points

Although the present study uses 85 monitoring sequences, the proposed framework is designed to operate under incomplete temporal conditions. The adaptive encoder architecture allows inference to be performed using shorter temporal sequences when only limited observations are available during early-stage monitoring. In such cases, the model can still generate probabilistic regime assignments, although classification confidence is reduced due to the lower temporal context available for latent-state estimation.

This capability is particularly relevant for operational infrastructure monitoring, where newly instrumented assets may initially contain only a small number of observations. Rather than requiring fully accumulated long-term time series, the framework supports progressive probabilistic assessment as new measurements become available over time.

The resulting probabilistic classifications for the 34 monitoring points are illustrated in Figure 2. A transitional tendency mode is assigned to points where the classification confidence remains below 90%, reflecting increased uncertainty or mixed behavioural characteristics between neighbouring regimes.

4.3 Discussion

The current results are obtained using only two modalities: displacement and rainfall data. Additional sensing modalities, such as vibration measurements, temperature records, and geometric metadata, may further improve classification robustness by providing complementary information not captured by the existing inputs alone.

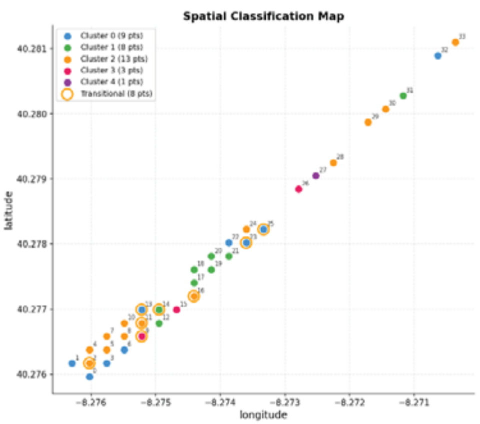

Fig. 2. Probabilistic regime classification results for 34 monitoring points on the studied geotechnical asset.

The Dirichlet process prior in the BGM automatically limits the number of active clusters, preventing overfitting when data are sparse or when modality information is limited. This property is critical for network-scale deployment where sensor availability varies across assets. Figure 3. point 32 exhibits a characteristic displacement trajectory consistent with active slope movement. The time series reveals an initial displacement phase followed by a period of stabilization around timestep 25–50, after which displacement accelerates progressively ending at 0.214 after fluctuating around a temporal mean displacement of 0.221.

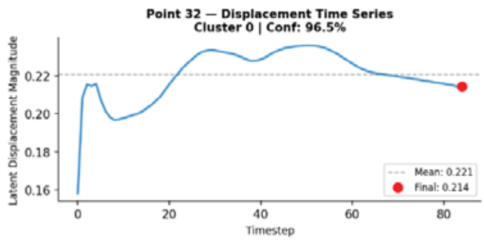

Fig. 3. Representative displacement time series for monitored Point 32.

The comparison against Cluster 0 (Figure. 4), the regime to which Point 32 is assigned, shows that the point's mean displacement (~0.22) and final displacement (~0.21) are generally consistent with the cluster profile (~0.25 and ~0.24, respectively), although slightly lower. Its volatility (~0.01) is also marginally lower than the cluster average (~0.02), indicating relatively stable temporal behaviour. Similarly, the trend value for Point 32 (~0.05) remains close to the cluster mean (~0.07), suggesting a moderate increasing displacement tendency consistent with the regime characteristics. The maximum and average maximum displacement values (~0.24 and ~0.22) are slightly below the cluster averages (~0.29 and ~0.24), indicating that Point 32 exhibits the same general deformation behaviour as Cluster 0 but at a somewhat lower intensity. Overall, the point closely follows the characteristic signature of Cluster 0, reflecting a stable and moderately evolving displacement pattern consistent with the assigned regime.

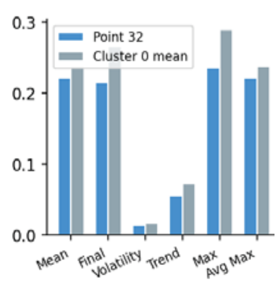

Fig. 4. Detailed regime comparison for monitored point 32.

5. Conclusions

This paper has presented a multi-modal physics-aware AI framework for unsupervised slope stability assessment, combining self-supervised latent encoding, FiLM-based environmental modulation, Kalman state estimation, and Bayesian Gaussian Mixture clustering. The framework was demonstrated on 671 monitoring points with InSAR horizontal displacement and precipitation data, identifying five dynamic response regimes with a silhouette score of 0.347.

The pipeline transforms raw heterogeneous observations into probabilistic behavioural regimes. Key contributions include: (i) a modular encoder architecture that explicitly models the conditional influence of rainfall on displacement through FiLM modulation; (ii) a vectorised Kalman filter with encoder-derived dynamic observation noise for uncertainty-aware state estimation; (iii) sliding-window preprocessing that enables near-real-time classification from limited temporal observations; and (iv) a Bayesian probabilistic inference stage that provides soft regime assignments for new monitoring points.

Future work will focus on the integration of vibration, strain, and temperature data from additional transportation infrastructure assets, including bridges and retaining structures, to evaluate the transferability and generalization capability of the proposed framework across different infrastructure typologies.

6. Acknowledgements

This work was funded by the ANI (Agência Nacional de Inovação) through the R&D project "WISE - Wireless Transportation Infrastructure Safety Evaluation", project number 14383, operation code COMPETE2030-FEDER-00574700, cofunded by the Innovation and Digital Transition Program (COMPETE 2030), by the PORTUGAL 2030 and by the European Union.
Also, this work is financed by national funds through FCT – Foundation for Science and Technology, under grant agreement CEECINST/00018/2021/CP2806/CT0006 attributed to the 2nd author. This work was also supported by FCT/MCTES under the R&D Unit Institute for Sustainability and Innovation in Structural Engineering (ISISE), under the references UID/4029/2025 (https://doi.org/10.54499/UID/04029/2025), UID/PRR/04029/2025 (https://doi.org/10.54499/UID/PRR/04029/2025) and UID/PRR2/04029/2025 (https://doi.org/10.54499/UID/PRR2/04029/2025), and under the Associate Laboratory Advanced Production and Intelligent Systems ARISE under reference LA/P/0112/2020 (https://doi.org/10.54499/LA/P/0112/2020).

References

[1] Casagli, N., Intrieri, E., Tofani, V., Gigli, G., Raspini, F., 2023. Landslide detection, monitoring and prediction with remote-sensing techniques. Nature Reviews Earth & Environment, 4, 51–64. https://doi.org/10.1038/s43017-022-00373-x
[2] Scaioni, M., Longoni, L., Melillo, V., Papini, M., 2014. Remote sensing for landslide investigations: An overview of recent achievements and perspectives. Remote Sensing, 6, 9600–9652. https://doi.org/10.3390/rs6109600
[3] Jiang, Y., Liao, M., Zhou, Z., Shi, X., Zhang, L., Bai, T., 2016. Landslide deformation analysis by coupling deformation time series from SAR data with hydrological factors through data assimilation. Remote Sensing, 8(3), 179. https://doi.org/10.3390/rs8030179
[4] Vaccari, A., Batabyal, T., Tabassum, N., Hoppe, E., Bruckno, B.S., Acton, S., Deaton, J., 2018. Integrating remote sensing data in decision support systems for transportation asset management. Transportation Research Record, 2672, 23–35.
[5] Casagli, N., Frodella, W., Morelli, S., et al., 2017. Spaceborne, UAV and ground-based remote sensing techniques for landslide mapping, monitoring and early warning. Geoenvironmental Disasters, 4, 1–23.
[6] Scaioni, M., Longoni, L., Melillo, V., Papini, M., 2014. Remote sensing for landslide investigations. Remote Sensing, 6, 9600–9652.
[7] Hormes, A., Adams, M., Amabile, A., et al., 2020. Innovative methods to monitor rock and mountain slope deformation. Geomechanics and Tunnelling, 13. https://doi.org/10.1002/geot.201900074
[8] Castagnetti, C., Bertacchini, E., Corsini, A., Rivola, R., 2014. A reliable methodology for

monitoring unstable slopes: The multi-platform and multi-sensor approach. Proceedings of SPIE, 9245.
[9] Vaccari, A. et al., 2018. Integrating remote sensing data in decision support systems for transportation asset management. Transportation Research Record, 2672, 23–35.
[10] Zhang, M., Wei, J., 2025. Analysis of slope stability based on four machine learning models: An example of 188 slopes. https://doi.org/10.3311/PPci.37630
[11] Bai, G., Hou, Y., Wang, B., et al., 2022. Performance evaluation and engineering verification of machine learning-based prediction models for slope stability. Applied Sciences, 12(15), 7890.
[12] Shu, T., Tao, W., Lu, H., et al., 2023. A hybrid random forest and least squares support vector machine model based on particle swarm optimisation for slope stability prediction. https://doi.org/10.1155/2023/6651323
[13] Tinoco, J., Correia, A.G., Cortez, P., Toll, D.G., 2019. Artificial neural networks for rock and soil cutting slopes stability condition prediction. https://doi.org/10.1007/978-3-030-02032-3_10
[14] Tinoco, J., Alberto, A., Oliveira, P.J.V., Lemos, L., Gomes Correia, A., 2019. Data-driven approach to predict unconfined compression strength. https://doi.org/10.32075/17ECSMGE-2019-0598
[15] Tao, W., 2013. Application of support vector machine for prediction of slope stability coefficient considering the influence of rainfall and water level. Applied Mechanics and Materials, 851, 840–845.
[16] Segalini, A., Carri, A., Valletta, A., Martino, M., 2019. Innovative monitoring tools and early warning systems for risk management. Geosciences, 9(2), 62.
[17] West, M., 2020. Bayesian forecasting of multivariate time series: scalability, structure uncertainty and decisions. https://doi.org/10.1007/s10463-019-00741-3
[18] von Rueden, L., Mayer, S., Sifa, R., Bauckhage, C., Garcke, J., 2020. Combining machine learning and simulation to a hybrid modelling approach: Current and future directions. https://doi.org/10.1007/978-3-030-44584-3_43
[19] Miao, P., Zhou, C., Wu, Y., Hu, W., Luo, D., 2025. Enhancing bridge performance assessment with a hybrid attention-based LSTM and hidden Markov model using sparse inspection data. Engineering Structures. https://doi.org/10.1016/j.istruc.2024.107995