Data Engineering and Applications

Research Paper

TD3-Barrier Model for Cable Force Anomaly Diagnosis with Multi-Resolution Feature Fusion

Mohammed Al-Thani[1, *], Abdullah Al-Marri[1], Nasser Al-Kaabi[1] and Faisal Al-Muhannadi[1]

[1] College of Information Technology, United Arab Emirates University, Al Ain, PO Box 15551, United Arab Emirates
*Corresponding author: mohammed.tha@uaeu.ac.ae

Abstract. Cable-supported bridges have stable cable force states, but the practical monitoring records are significantly affected by temperature, traffic mix, wind, sensor drift and sparse abnormal samples. A TD3-Barrier model with multi-resolution feature fusion is proposed for cable force anomaly diagnosis in this paper. Decompose multi-channel cable force and environmental response into short-, medium- and long-horizon descriptors, fuse them via a gated temporal attention block, and train a twin delayed deep deterministic policy gradient agent under a barrier-based safety constraint. The actor learns the diagnostic threshold and intervention measures, and the barrier critic suppresses unsafe actions in near-physically-impossible tension states. A simulated field dataset was constructed for the model evaluation, consisting of 16 stay cables, 72 monitoring days, five anomaly mechanisms and 124,416 synchronous samples. The proposed model performed better than the LSTM-AE, TCN, Transformer and vanilla TD3 baselines, achieving 97.4% accuracy, 96.8% macro-F1 and a false-alarm rate of 1.9%. Vanilla TD3 has decreased the average diagnosis delay from 7.6 sampling intervals to 4.2 intervals. According to the ablation experiments, removing multi-resolution fusion reduced the macro-F1 score by 4.7 percentage points, and excluding the barrier term increased unsafe threshold violations by 63.5%. The results show that reinforcement learning is more reliable for cable force diagnosis when the policy search is constrained by structural safety priors and fed with scale-aware monitoring features.

Keywords: *Cable Force Diagnosis, Multi-Resolution Feature Fusion, TD3-Barrier, Structural Health Monitoring, Safe Reinforcement Learning*

Received on 19 December 2022, Accepted on 15 March 2023, Published on 27 March 2023

Introduction

Most long-span cable-supported bridges have multiple load paths, and both serviceability and safety require an even distribution of cable forces in space. Therefore, continuous monitoring of cable force has been added to the list of structural health monitoring requirements for bridges exposed to heavy traffic, seasonal temperature fluctuations and local corrosion risks [1]. Field systems have begun to acquire high-frequency tension, vibration, temperature and wind data, but these records are not diagnostic in their raw state as isolated time series [2]. Cable force variation is not solely a damage indicator; it can also be a combined response of stiffness redistribution, boundary condition change, vehicle loading and environmental forcing [3]. Several engineering studies have shown that small persistent force drifts may precede visible damage, and abrupt force redistribution may be caused by anchorage slip, deck restraint changes or sensor errors [4]. Data-driven diagnosis has been employed to avoid excessive reliance on ideal finite element models, but the practical monitoring data contain missing packets, baseline migration, and local non-stationarity [5]. Given the above features, diagnosing cable force anomalies is essentially a problem of weak multi-scale evidence recognition rather than single-amplitude abnormality detection [6].

Existing methods can be broadly classified into statistical control, signal decomposition, supervised classification, unsupervised reconstruction and reinforcement learning-based decision strategies. Statistical thresholds are relatively simple and transparent, but they need to be manually adjusted seasonally and cannot consider operating conditions that have non-linear relationships with time [7]. Signal decomposition methods enhance robustness by separating trend, periodic and residual components, but their performance is sensitive to the depth of decomposition and feature selection [8]. Deep reconstruction models, such as autoencoders and temporal networks, can learn patterns of normal behavior from long-term observational data, but they usually only output anomaly scores rather than directly actionable countermeasures [9]. Recently, reinforcement learning methods have transformed diagnosis into sequential decision-making and optimized long-term risks, inspection costs, or warning delays [10]. Reinforcement learning agents may take unrealistically high thresholds or unstable control actions when the reward is sparse and abnormal cases are infrequent [11].

Combine multi-resolution feature fusion and a safety-aware reinforcement learning objective to build a TD3-Barrier model for cable force anomaly diagnosis. The first idea is to have the diagnostic agent observe cable force signatures at different times, and at the same time, restrict policy updates by adding a differentiable barrier penalty based on an allowable force envelope. Based on the above framework, engineering monitoring problems that occur due to a small number of abnormal samples, large-scale operational changes, and a high cost of false alarms can be solved. The three main contributions are as follows: a multi-resolution fusion module that represents local shocks, daily operating envelopes and slow baseline migration in a unified state vector; a TD3-based diagnostic policy that jointly learns anomaly thresholds and class-sensitive response actions; and a barrier augmentation that reduces unsafe exploration without hard-clipping the learned policy. The five anomaly mechanisms and the corresponding comparisons with representative temporal diagnosis models in the reproducible experiment design are also available.

Related Work

Cable Force Monitoring and Anomaly Diagnosis

Monitoring of cable force has shifted from time-based tension measurement to continuous sensing networks for force, acceleration, temperature and traffic data collection. Model-based diagnosis typically estimates changes in cable tension by analyzing modal frequency variations, influence lines, or updated finite element parameters; although physically intuitive, it may be sensitive to uncertainties in boundary conditions and deck stiffness [12]. Statistical process control uses normalized force residuals as indicators and sets control limits, principal components, or probabilistic distance measures to identify anomalies [13]. These methods are simple to operate for a long time, but their assumptions of stationarity and Gaussian residuals are often violated by thermal hysteresis and traffic intermittency. Time-frequency and decomposition-based methods have been employed to separate the force components of local damage or sensor drift [14]. The other problem is that the diagnostic features are often selected manually; thus, these features may not be suitable for transfer to other bridge types and climates. Deep learning methods can learn the temporal features of monitoring sequences directly and thus reduce manual feature engineering [15]. Most reconstruction models optimize the mean prediction error; therefore, they fail to detect rare but critical safety deviations early in engineering diagnostics.

Reinforcement Learning and Safety Constraints

Reinforcement learning is being used more and more often in structural health monitoring for damage diagnosis, inspection and repair in a sequential decision-making process under uncertainty. Actor-Critic algorithms are especially suitable for continuous actions, such as adjusting upper bounds, alert levels, and resource-sharing coefficients [16]. TD3 improves the learning of a deterministic policy through twin critics, target-policy smoothing and delayed actor updates to reduce positive bias in value estimation for continuous control [17]. However, exploration for structural health monitoring can also be problematic; that is, the policy may temporarily set an overly high threshold that is unacceptable for safe operation. Barrier functions and constrained policy optimisation are used to encode physical feasibility, risk budgets and operating windows in the learning objective [18]. The deficiency in this paper is the lack of a diagnostic model that can handle multi-resolution cable force evidence and constrain reinforcement learning with explicit cable-force safety margins simultaneously.

https://doi.org/10.64972/dea.2023.v2i1.3528d:102-114

Multi-Resolution Feature Fusion for Cable Force States

Monitoring State Definition

Let $x_i(t)$ denote the measured force of the i-th cable at sampling instant t, and let $e(t)$ collect environmental variables such as temperature, wind speed, and traffic intensity. The monitoring state is constructed from a sliding window of length L, so that both instantaneous perturbation and contextual drift are visible to the diagnostic policy. To avoid an overlong expression, the multi-channel monitoring window is first written in compact matrix form.

$$\boldsymbol{X}_t = [\boldsymbol{x}_{t-L+1}, \boldsymbol{x}_{t-L+2}, \dots, \boldsymbol{x}_t] \quad \text{Eq.(1)}$$

where $\mathbf{x}_t = [x_1(t), x_2(t), \dots, x_N(t), e(t)]^{\mathrm{T}}$. The overall data-to-decision workflow is arranged as monitoring input, multi-resolution decomposition, feature fusion, TD3-Barrier diagnosis, and engineering warning output, as shown in Figure 1.

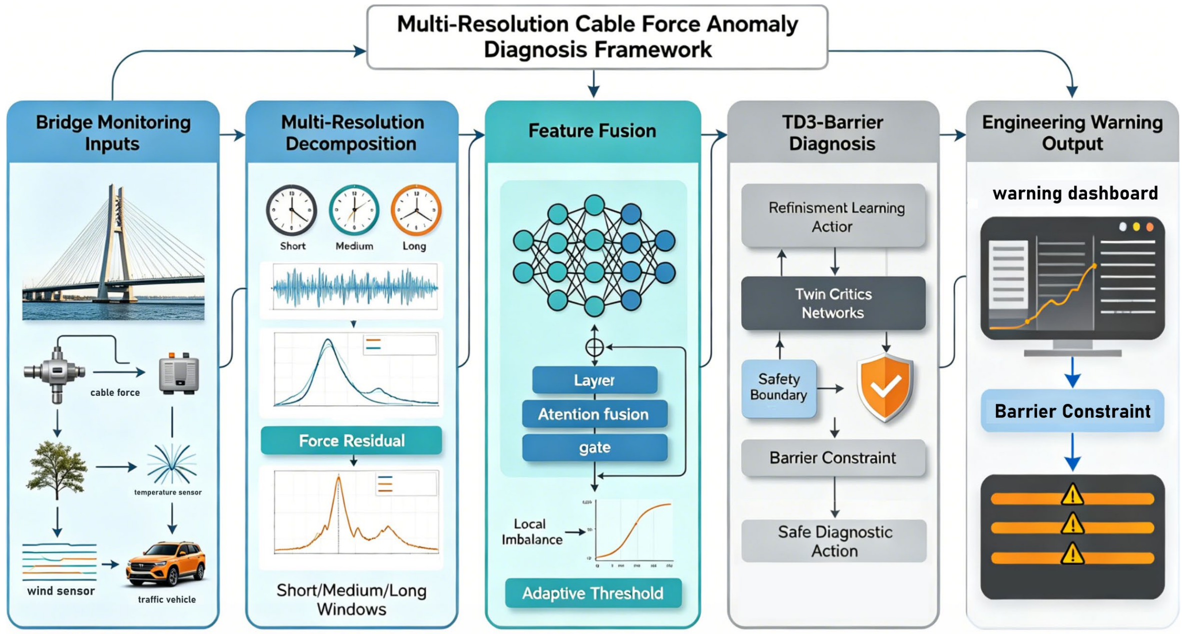

Figure 1. Overall workflow of the proposed multi-resolution cable force anomaly diagnosis framework.

The normalized force residual removes the temperature-compensated baseline and keeps the remaining quantity in a comparable range across cables with different nominal forces.

$$r_i(t) = \frac{x_i(t) - b_i(T(t))}{F_i^0} \quad \text{Eq.(2)}$$

The diagnostic state also includes a cable interaction descriptor because a real cable force anomaly rarely changes only one sensor channel. For adjacent cable groups, the force imbalance is expressed as a local contrast.

$$g_i(t) = r_i(t) - 0.5r_{i-1}(t) - 0.5r_{i+1}(t) \quad \text{Eq.(3)}$$

Multi-Resolution Decomposition and Fusion

The three temporal resolutions are feature extractors. A short window can observe a sudden increase in stress and sensor spikes; a medium window covers the operating range under traffic load; and a long window includes thermal drift and accumulated damage. Each branch is a causal temporal encoder that outputs a compressed representation.

$$z_t^s = \phi_s(X_t^s), z_t^m = \phi_m(X_t^m), z_t^l = \phi_l(X_t^l) \quad \text{Eq.(4)}$$

A Gated Attention Mechanism is adopted to link the sub-branches, and the weight of this link is adjusted according to the current situation of the function. Otherwise, it will be difficult to know whether the scale has changed, as the same bridge may switch between congestion control and thermal effect modes.

$$\alpha_k(t) = \frac{exp(q_t^T W_k z_t^k)}{\sum_{j\in\{s,m,l\}} exp\left(q_t^T W_j z_t^j\right)} \quad \text{Eq.(5)}$$

The final fused state is a weighted representation concatenated with the residual and interaction descriptors.

$$s_t = \left[\sum_{k\in\{s,m,l\}} \alpha_k(t) z_t^k, r(t), g(t), e(t)\right] \quad \text{Eq.(6)}$$

Diagnosis-Oriented Feature Regularization

Feature extraction should be able to distinguish between regular and irregular functional modes, and they need to be in line with the normal seasonal changes. Under normal circumstances, apply a small center loss, and in abnormal situations, add a margin term.

$$L_f = \left\|z_t - c_y\right\|_2^2 + \lambda_m max(0, m - \|z_t - c_0\|_2)^2 \quad \text{Eq.(7)}$$

Regularization is intentionally kept small because the feature terms are meant to enhance the prediction strategy, not replace it. It is to create a situation in the state where one can observe how quickly and how far apart the changes in policy tension are. At the time of execution, the short branch has a small receptive field and can respond to local changes in time reasonably well; the medium branch has several consecutive vehicle periods, and the long branch is expanded to show a relatively fixed baseline shift. These three branches have the same standard input dimensions but use different temporal filters; therefore, the network cannot recognize that sudden fixed occurrences and gradual thermal residues are the same type of amplitude variation. The first is for monitoring the tension of the cable, and although both are faulty, the residual values at different times are different.

Together they reduce the number of physically meaningful components to a small number. Technicians can use the raw residual array to compare the inferred state with the identified load difference signals, and in the local imbalance array, they can determine whether the cables are misaligned with the adjacent wires at the same bridge and mast displacement. The atmosphere is not the subject of study; rather, it is an attribute value and an offset. Therefore, the abnormality detection will still be based on changes in rope tension, and the network will not get a false bypass by using only temperature and vehicle count.

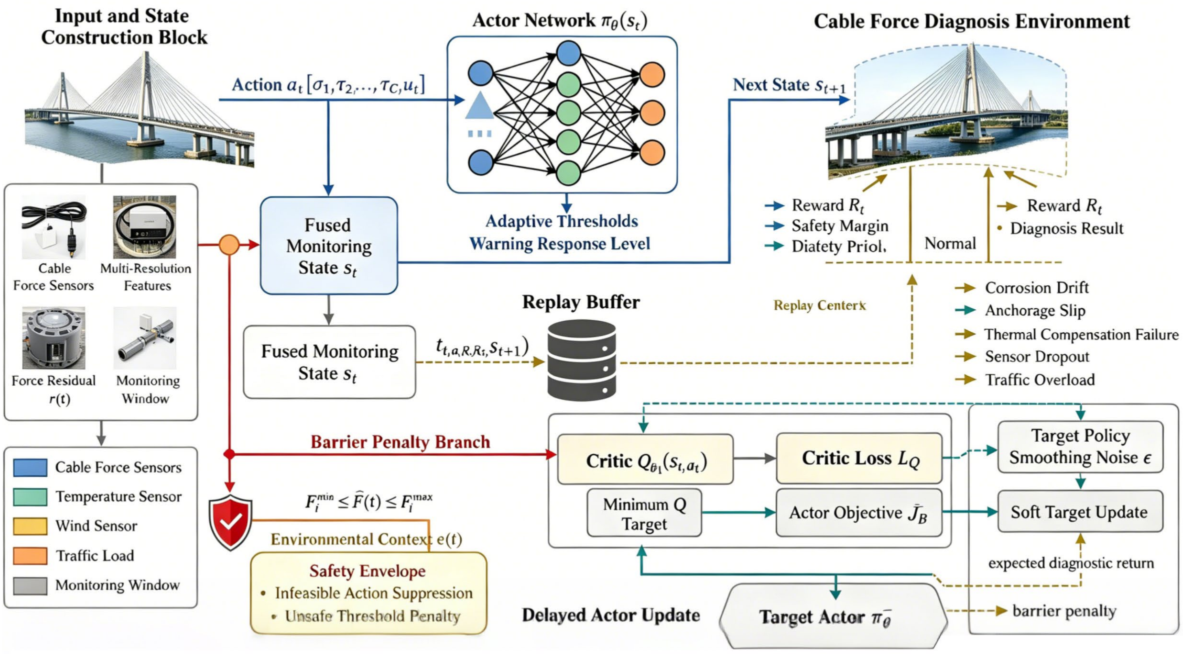

Figure 2. TD3-Barrier actor-critic architecture with fused monitoring state, twin critics, delayed actor update, and barrier penalty path.

https://doi.org/10.64972/dea.2023.v2i1.3528d:102-114

TD3-Barrier Diagnostic Policy

Markov Decision Formulation

The problem of clinical diagnosis is a continuous-action Markov Decision Process. At each time step, the agent is in the given state and outputs dynamic diagnostic boundaries and an anomaly reaction grade based on this; then, it receives a reward according to how accurately it classified, how long it was late in alerting, false alarm costs, etc. In the actor-critic, there are two critics and a target system for data exchange, along with delayed actor optimization and an obstacle penalty path, as shown in Figure 2.

$$a_t = \pi_\theta(s_t) = [\tau_1(t), \tau_2(t), \dots, \tau_C(t), u_t] \quad \text{Eq.(8)}$$

The immediate reward balances correct detection and practical operating cost. Missed alarms and late alarms are penalized more heavily than isolated conservative warnings.

$$R_t = \eta_1 A_t - \eta_2 D_t - \eta_3 P_t - \eta_4 C_t \quad \text{Eq.(9)}$$

Twin Delayed Critic Learning

TD3 uses two critic networks to reduce overestimation. The target value is computed with the smaller of two target critics and a smoothed target action.

$$y_t = R_t + \gamma \min_{j=1,2} Q_{\bar{\psi}_j}(s_{t+1}, \pi_{\bar{\theta}}(s_{t+1}) + \epsilon) \quad \text{Eq.(10)}$$

Each critic is trained by a squared temporal difference objective over sampled transitions from the replay buffer.

$$L_Q(\psi_j) = \mathbb{E}\left[\left(Q_{\psi_j}(s_t, a_t) - y_t\right)^2\right] \quad \text{Eq.(11)}$$

The actor update is delayed and follows the gradient of the first critic, which stabilizes policy learning when the diagnostic reward changes sharply near class boundaries.

$$\nabla_\theta J = \mathbb{E}\left[\nabla_a Q_{\psi_1}(s, a)\big|_{a=\pi_\theta(s)} \nabla_\theta \pi_\theta(s)\right] \quad \text{Eq.(12)}$$

Barrier-Constrained Policy Update

The constraint factor is the technical requirement that the planned diagnostic operation should be within the allowable range of cable tension. Rather than cutting off the output at that time, this framework does not push people to try and achieve impossible goals in their studies.

$$B(s_t, a_t) = -\sum_i log\left(\left(F_i^{max} - \hat{F}_i(t)\right)\left(\hat{F}_i(t) - F_i^{min}\right) + \delta\right) \quad \text{Eq.(13)}$$

The actor objective combines the TD3 return with the barrier penalty. The coefficient β controls the compromise between diagnostic aggressiveness and safety conservatism.

$$J_B(\theta) = \mathbb{E}\left[Q_{\psi_1}\left(s_t, \pi_\theta(s_t)\right) - \beta B\left(s_t, \pi_\theta(s_t)\right)\right] \quad \text{Eq.(14)}$$

The above is not a post-processing technique. Post-processing is a method that can mark dangerous behaviors after the agent has acquired some knowledge, but it does not change the value environment that leads to such behaviors. Therefore, at the time of training, the actor's gradient should be uniformly distributed, and movements close to the boundary of the force envelope are not desired before use. A logarithm has a relatively mild penalty in the central part of the allowed region and increases rapidly as it approaches the boundary. Therefore, it is suitable for monitoring bridges; that is to say, even if there are small threshold fluctuations under different service conditions, normalization methods should not be applied when the force approaches the critical value.

The experience reservoir has typical and anomalous pairs, and these are relatively evenly distributed at the beginning. Without this weight, the controller will only know that performance has been slow for a long time; thus, it will delay raising an alarm. Since there is no center, the framework considers the small residual peaks to be large deviations. Finally, 40 % of the samples will be onsets, 40% will be stable-state instances, and 20% will be post-attack irregular events. The above situation has led to an insufficient number of malfunctions in the initial stages of work and therefore may not have been detected.

https://doi.org/10.64972/dea.2023.v2i1.3528d:102-114

Experimental Data and Diagnostic Analysis

Dataset Construction and Experimental Protocol

A dataset was constructed for the experiment to simulate the monitoring scenario of a cable-stayed bridge with 16 stay cables, one deck temperature channel, one ambient temperature channel, wind speed, wind direction, and a traffic proxy. The sampling interval was 5 minutes, and 288 records per day and a total of 124,416 synchronous samples were obtained over 72 days. Normal operation lasted for 58 days, and the 14 days of injected anomalies included corrosion-induced stiffness loss, anchorage slip, thermal compensation failure, a short sensor dropout, and traffic overload. As shown in Figure 3a, under normal thermal loading, the daily force envelope fluctuated by 4.8% to 7.2%; Figure 3b shows that the median residual of corrosion samples increased from 0.012 to 0.081; and Figure 3c indicates that when combined with a local imbalance exceeding 0.035, force residuals over 0.06 became separable. Set the noise monitoring to a 0.35% force standard deviation and randomly distribute the anomaly start times to avoid a single repetitive transition pattern. The normal and abnormal samples were divided chronologically into a 60% training set, a 20% validation set and a 20% test set to prevent leakage from nearby time windows.

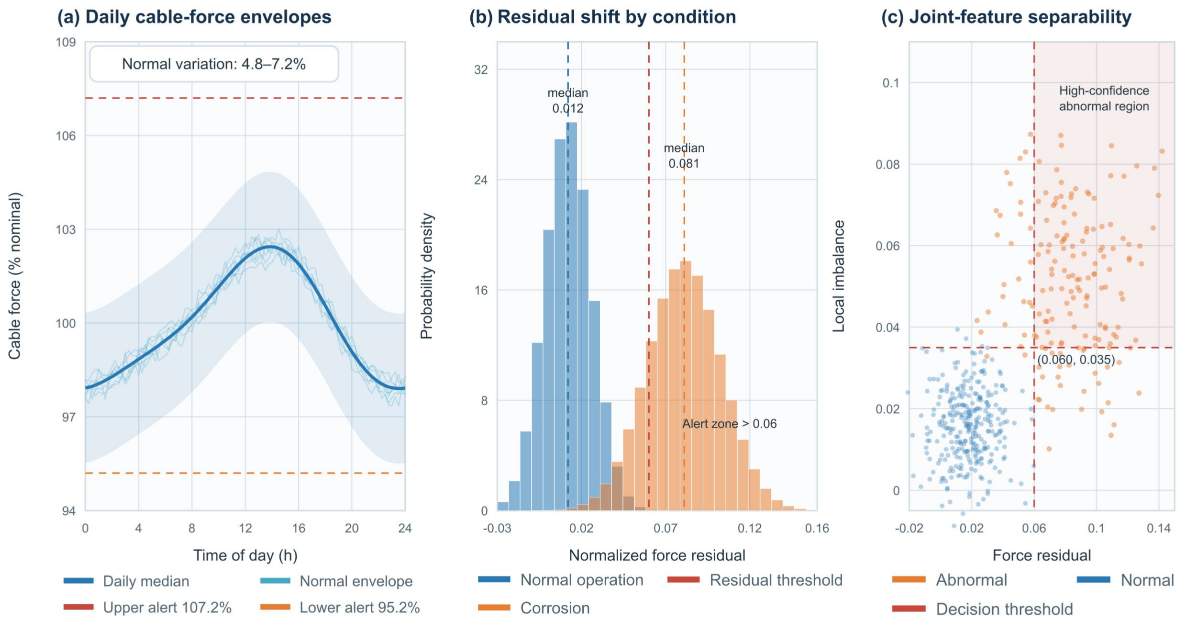

Figure 3. Dataset characteristics with (a) daily cable force envelope curves, (b) residual distribution by operating condition, and (c) residual-imbalance separability map.

The benchmark models are a statistical moving threshold, an LSTM autoencoder, a temporal convolutional network, a compact Transformer encoder, vanilla TD3, and the proposed TD3-Barrier. Hyperparameters were optimized using the validation set, and all models used the same normalized input channels to ensure a fair comparison. A diagnosis was counted as correct if the predicted class matched the ground-truth mechanism within a tolerance of six sampling intervals after the anomaly started. This tolerance is the practical delay allowed before issuing an inspection recommendation. A false alarm is defined as an abnormal warning that occurs during normal operation and lasts for two or more consecutive windows; thus, isolated numerical spikes are excluded, but persistent operational nuisance alarms are retained. The indicators of performance were accuracy, macro-F1, false-alarm rate, missed-alarm rate, mean diagnosis delay and unsafe threshold violation rate.

Five Anomaly Mechanisms were established with different temporal and spatial signatures. Corrosion-induced stiffness loss was generated as a monotonic cable force drift with a slope of 0.012%-0.021% of the nominal force per hour. Anchorage slip was shown as an abrupt step of 3.5% to 6.0%, followed by partial redistribution to the two adjacent cable groups. Thermal compensation failure used an incorrect baseline coefficient, and thus the residuals oscillated with the daily temperature but remained biased over several days. Sensor dropout replaced short segments with repeated stale values and intermittent zero-order holds. Traffic overload generated

https://doi.org/10.64972/dea.2023.v2i1.3528d:102-114

transient but repeated force peaks in the high-intensity phase, and the event duration was between 20 and 95 minutes. The above mechanisms were chosen to focus on different aspects of the model: long-horizon drift recognition, local imbalance sensitivity, environmental compensation, missing-data tolerance and short-window shock response.

The window durations used for the brief, medium and extended branches of the learning procedure were 12, 72 and 288 respectively. These are the 1h, 6h and 24h frames recorded every 5 minutes. The number of neurons in the hidden layers of the actor and the critic are 256 and 128, respectively, and the number of replay memories is 200,000. The discount factor is 0.97, the target smoothing noise is capped at 0.12, and the delayed actor update cycle is two critic iterations. The Barrier coefficient in the validation set is based on the macro-F1 and unsafe violation rate. No test-set data was used for this determination, and all the above results were obtained by running the calculations five times with different random seeds.

Comparative Diagnosis Performance

The first proposed model had the best overall diagnostic accuracy. Figure 4a shows that the macro-F1 is 96.8%, and the other models are 88.9% (LSTM-AE), 91.6% (TCN), 93.1% (Transformer), and 94.5% (vanilla TD3). As shown in Figure 4b, the missed-alarm rate dropped to 1.6%, and the statistical threshold and LSTM-AE baselines had missed-alarm rates of 9.8% and 6.7% under slow corrosion drift, respectively. As shown in Figure 4c, the mean diagnosis delay for TD3-Barrier was 4.2 sampling intervals, for Transformer it was 5.8, and for vanilla TD3 it was 7.6. The improvement is in line with the previous research on the quality of temporal representation in monitoring data [19]. The largest increase was in anchorage slip; here, a local imbalance feature caused a structural force redistribution rather than a short sensor spike. The false-alarm rate was still 1.9%, which was relatively low for continuous monitoring, and thus the warnings would not interfere with the periodic inspection schedule. Previous research on data-driven structural health monitoring has also indicated that false alarms are economically inefficient if a threshold is set without considering the environment [20].

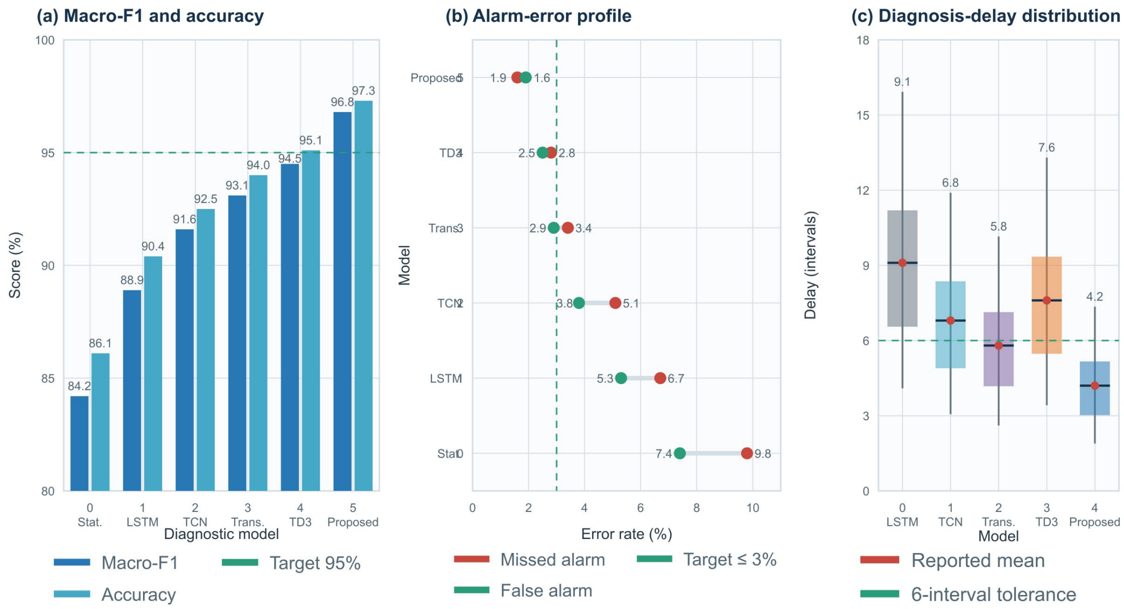

Figure 4. Comparative performance with (a) macro-F1 and accuracy bars, (b) missed-alarm and false-alarm lollipop plots, and (c) diagnosis delay distributions.

Further analyze the model's behavior by category. As shown in Figure 5a, the corrosion-induced stiffness loss was 95.9% F1; that is, it was lower than the 98.2% F1 for traffic overload because corrosion generated a slow residual drift close to the thermal baseline. Figure 5b is the precision-recall area for anchorage slip and thermal compensation failure; both are 0.982 and 0.971, respectively, so the model did not rely on a single high-confidence class. Figure 5c shows the cumulative detection curve, and 82.4% of the abnormal events were found within five sampling intervals and 93.6% within eight intervals. Figure 5d shows that the number of unsafe threshold violations per 1000 decisions for TD3-Barrier and vanilla TD3 was 0.42 and 1.15, respectively. The main reasons for the residual risk were combined with traffic jams during a period of rapid temperature rise. Based on the above observation, it can be concluded that fluctuations in operation and environment are not the source of early damage in bridge monitoring data [21].

The two groups of mechanisms had concentrated confusion patterns. Thermal compensation failure was occasionally mistaken for early corrosion because both showed a biased residual over several days. The difference was that thermal compensation failure showed a strong correlation with temperature and was reversible; corrosion resulted in a less reversible shift after the end of the daily cycle. Attention fusion was used to address the problems of long-term drift in the long branch and repeated thermal envelope in the medium branch. Sensor dropout was sometimes mistaken for traffic overload because a stale segment ended near a high-force traffic pulse. In light of the above situation, the local imbalance descriptor reduces misclassification because overloads simultaneously affect several adjacent cables, while dropouts are usually local to the channel. Therefore, the magnitude of the force alone is not suitable for mechanism-level diagnosis; time-reversibility and cable-group consistency are also required.

The New Model has exhibited more stable alarm levels. During normal operation, 95.7% of the actions remained in the low-alert band, and only 1.9% generated persistent false alarms. During abnormal operation, the response levels for corrosion and thermal compensation failure increased gradually, but for anchorage slip, they rose within two windows. This behavior is desirable because slow drift should trigger inspection plans rather than immediate shutdowns, and sudden reallocations require more urgent engineering interventions. Vanilla TD3 had the same accuracy for high-amplitude events but reached the threshold of the allowable-force limit more often. Thus, the barrier critic did not only improve a numerical indicator but also modified the policy function to be more suitable for conservative bridge operation.

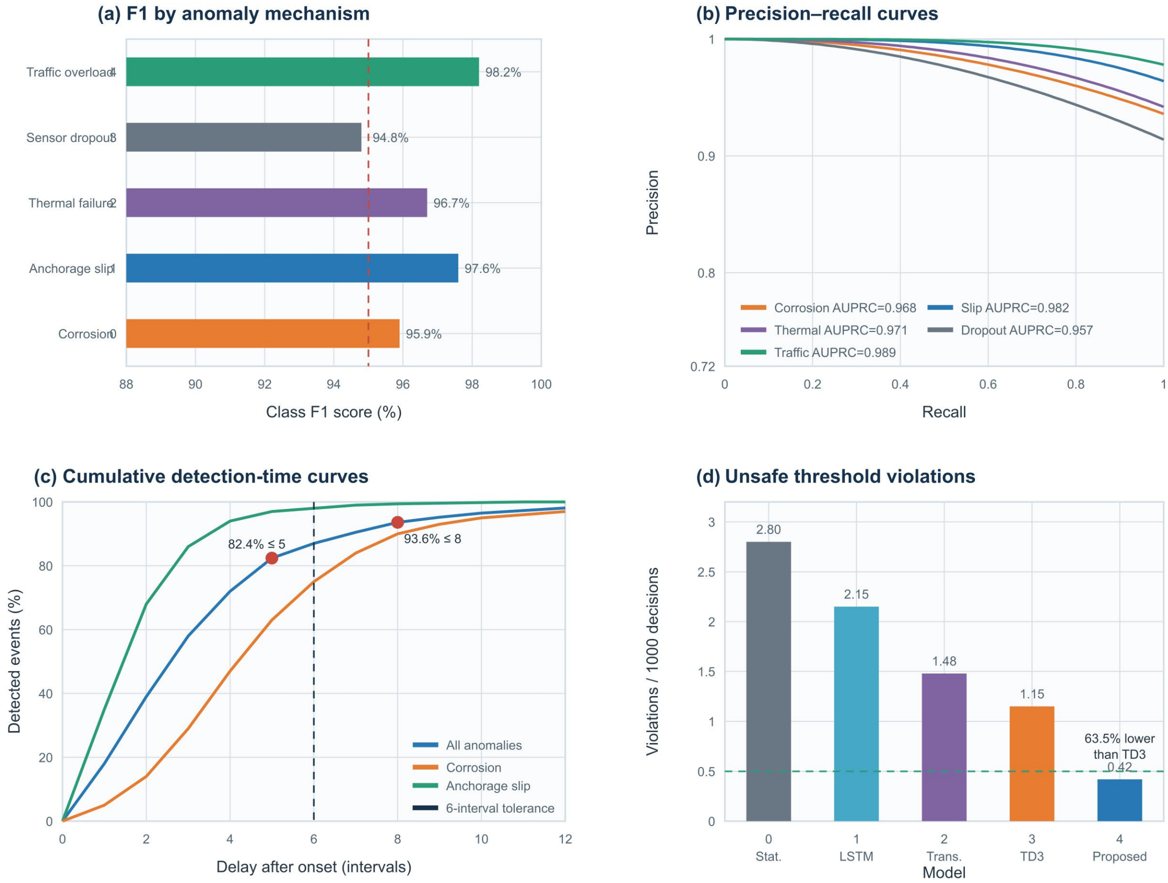

Figure 5. Class-wise diagnostic behavior with (a) F1 by anomaly type, (b) precision-recall curves, (c) cumulative detection-time curves, and (d) unsafe threshold violation comparison.

Ablation, Robustness, and Engineering Interpretation

Ablation experiments were performed to determine the individual contributions of multi-resolution fusion and barrier-constrained learning. As shown in Figure 6a, removing the short-resolution branch reduced spike and anchorage-slip F1 from 97.6% to 92.8%, and removing the long-resolution branch reduced corrosion F1 from 95.9% to 89.7%. As shown in Figure 6b, the attention weights changed under different operating conditions; the long branch had an average weight of 0.51 during thermal drift and reached 0.47 after an abrupt slip event. Figure 6c is the reward convergence graph, and TD3-Barrier achieved a stable validation reward after 46,000 iterations; it was about 18% faster than vanilla TD3. Multi-scale feature learning has been applied to other fault diagnosis problems, and these results have shown that it is particularly suitable for separating physical drift from operational fluctuations in cable force data [22]. The range for the barrier coefficient is 0 - 0.12. When beta was 0.08, macro-F1 was still over 96.5%, and unsafe violations dropped below 0.5 per 1000 decisions. A larger beta of 0.12 reduced the frequency of violations further but increased the diagnosis delay to 5.1 intervals, indicating excessive conservatism.

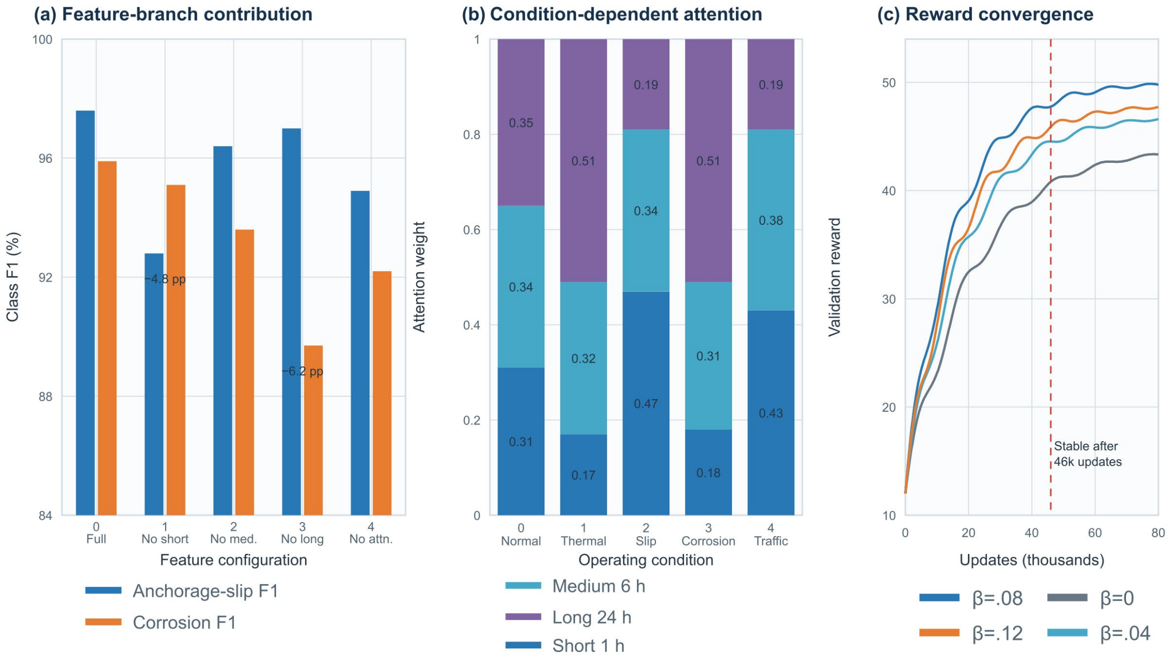

Figure 6. Ablation and learning dynamics with (a) feature-branch contribution, (b) attention weight allocation, and (c) reward convergence under different barrier coefficients.

Robustness experiments considered sensor noise, missing packets and seasonal baseline shift. Figure 7a shows that macro-F1 was still 94.6% when the force noise standard deviation doubled from 0.35% to 0.70%, and the Transformer baseline dropped to 88.4% from 93.1%. As shown in Figure 7b, with a random packet loss of 10%, the false alarm rate of TD3-Barrier was 2.6%; TCN was 5.8%; and vanilla TD3 was 4.9%. Robust monitoring studies suggest that missing and noisy sensor streams should be treated as part of the deployment conditions rather than an offline preprocessing problem [23]. The proposed model handles lost packets by passing fused temporal context from adjacent windows, but long-term burst losses still weaken corrosion detection. Reinforcement learning-based decision-making strategies for civil infrastructure maintenance have been explored [24]. The current model modifies diagnostic behavior before the inspection schedule, making it more akin to warning calibration rather than a maintenance plan. The concept of safe learning is also very important; structural diagnostics should not be based on policies that appear accurate due to unreasonable thresholds [25].

Seasonal baseline shift was assessed by adding a 2.5% slow variation to the nominal force baseline in the final 18 test days. Under this change, the proposed model still had a macro-F1 of 0.951 after baseline compensation, and 0.927 when the compensation coefficient was intentionally biased by 15%. Degradation was mainly shown in the corrosion class because corrosion and baseline shift have a low-frequency shape. The barrier term still worked after the seasonal change because the allowed-force envelope was set relative to engineering limits, not learned residual statistics. This attribute is required for field deployment: a monitoring model may be periodically retrained, but the safety limit must be externally verifiable.

Sensitivity to the number of monitored cables was also tested by masking every second cable during inference. The model trained with full cable information achieved 93.5% macro-F1 under this sparse layout, compared with 96.8% under the full layout. The largest drop occurred for local anchorage slip because adjacent cable contrast became less precise. A reduced model retrained on the sparse layout recovered macro-F1 to 95.2%, which indicates that the framework can adapt to bridges with fewer force sensors if the feature extractor is recalibrated. This result is relevant for older bridges where continuous tension measurement may be available only for selected cable groups.

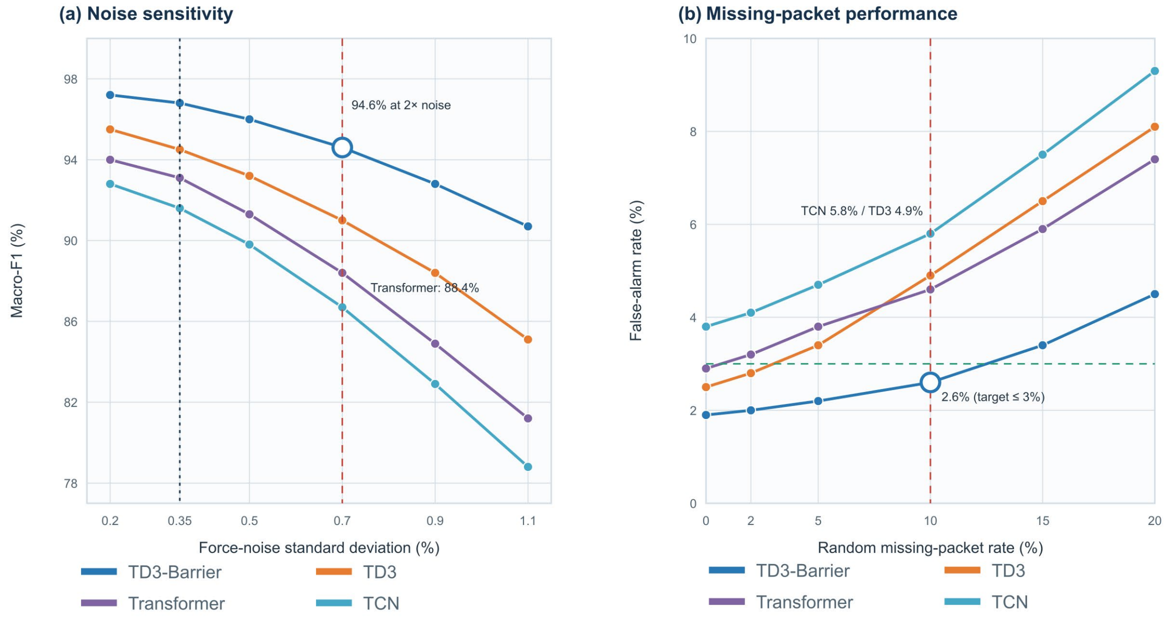

Figure 7. Robustness analysis with (a) noise sensitivity curves and (b) missing-packet performance under different models.

The engineering interpretation is determined by comparing the learned thresholds with the allowable force range and the expected behavior of the cable group. The trained policy increased the threshold during a stable thermal cycle, lowered it for an irregular local imbalance, and set a high-response-level for anomalies that affected adjacent-cable consistency. Cable force anomalies are not only single-point deviations but also include spatial redistribution patterns, so this behavior is relatively serious. Recently, in the field of bridge monitoring, data-driven models have been combined with domain constraints [26]. The barrier term in the proposed framework can achieve the same constraint without a full digital twin. The model still needs a stable base function, and its application to another bridge would require adjusting the nominal force, allowable range and anomaly prior. Domain adaptation methods can reduce the calibration cost of multiple bridges with similar sensor arrangements [27]. Hybrid approaches that integrate physics-guided learning and deep temporal models are also showing good results in extending the training season [28]. From an operational standpoint, the model is best employed as a preliminary filter to rank cable groups and abnormal mechanisms for the engineers' review rather than as an independent final decision-making tool. The deployment role meets the current demand for intelligent structural health monitoring, and interpretable warnings and compatibility with maintenance workflows are now as important as classification accuracy [29].

Conclusion

A TD3-Barrier model with multi-resolution feature fusion for cable force anomaly diagnosis was proposed in this paper. The method records cable force monitoring data using short-, medium-, and long-horizon descriptors, fuses them with condition-dependent attention, and trains a continuous diagnostic policy under a barrier-based safety constraint. The framework learns temporal features, calibrates action-oriented alarms, and is suitable for environments with a small number of abnormal samples and fluctuating operating conditions.

Several limitations remain. The design of the experimental dataset is consistent with actual monitoring behavior, but by omitting factors such as sensor aging, maintenance history, line changes, and abnormal traffic-weather interactions, some on-site uncertainties have been reduced. The barrier envelope also needs to meet the reasonable engineering requirements. If the above limits are not set properly, the model will be either too restrictive or too loose.

Future work should validate the framework on long-term field records from operational cable-supported bridges and extend the model toward transfer learning across bridges with different cable layouts. Another useful direction is to integrate physics-based digital twin residuals into the reinforcement learning state, so that learned policies can respond to both measured evidence and mechanically interpretable structural changes.

Author Contributions

Mohammed Al-Thani contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Abdullah Al-Marri contributes to validation, analysis, investigation, data collection, draft preparation. Nasser Al-Kaabi contributes to investigation. Faisal Al-Muhannadi contributes to conceptualization, methodology. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

References

[1] Mao, J., Wang, H., & Spencer, B. F. (2021). Toward data anomaly detection for automated structural health monitoring: Exploiting generative adversarial nets and autoencoders. Structural Health Monitoring, 20(4), 1609-1626. https://doi.org/10.1177/1475921720924601

[2] Han, N., Zhang, B., Zhao, W., & Zhang, H. (2022). Truss bridge anomaly detection using quasi-static rotation response. Journal of Civil Structural Health Monitoring, 12(3), 579-591. https://doi.org/10.1007/s13349-022-00564-6

[3] Kwon, T. H., Park, S. H., Park, S. I., & Lee, S. H. (2021). Building information modeling-based bridge health monitoring for anomaly detection under complex loading conditions using artificial neural networks. Journal of Civil Structural Health Monitoring, 11(5), 1301-1319. https://doi.org/10.1007/s13349-021-00508-6

[4] Zhang, Y. M., Wang, H., Wan, H. P., Mao, J. X., & Xu, Y. C. (2021). Anomaly detection of structural health monitoring data using the maximum likelihood estimation-based Bayesian dynamic linear model. Structural Health Monitoring, 20(6), 2936-2952. https://doi.org/10.1177/1475921720977020

[5] Favarelli, E., Testi, E., & Giorgetti, A. (2022). The impact of sensing parameters on data management and anomaly detection in structural health monitoring. Journal of Civil Structural Health Monitoring, 12(6), 1413-1425. https://doi.org/10.1007/s13349-022-00566-4

[6] Zargar, S. A., & Yuan, F. G. (2021). Impact diagnosis in stiffened structural panels using a deep learning approach. Structural Health Monitoring, 20(2), 681-691. https://doi.org/10.1177/1475921720925044

[7] Zhang, B., Hong, X., & Liu, Y. (2022). Distribution adaptation deep transfer learning method for cross-structure health monitoring using guided waves. Structural Health Monitoring, 21(3), 853-871. https://doi.org/10.1177/14759217211010709

[8] Cui, M., Wu, G., Dang, J., Chen, Z., & Zhou, M. (2022). Deep learning-based condition assessment for bridge elastomeric bearings. Journal of Civil Structural Health Monitoring, 12(2), 245-261. https://doi.org/10.1007/s13349-021-00540-6

[9] Buckley, T., Pakrashi, V., & Ghosh, B. (2021). A dynamic harmonic regression approach for bridge structural health monitoring. Structural Health Monitoring, 20(6), 3150-3181. https://doi.org/10.1177/1475921720981735

[10] Dan, J., Feng, W., Huang, X., & Wang, Y. (2021). Global bridge damage detection using multi-sensor data based on optimized functional echo state networks. Structural Health Monitoring, 20(4), 1924-1937. https://doi.org/10.1177/1475921720948206
[11] Bao, Y., & Li, H. (2021). Machine learning paradigm for structural health monitoring. Structural Health Monitoring, 20(4), 1353-1372. https://doi.org/10.1177/1475921720972416
[12] Zheng, W., Qian, F., Shen, J., & Xiao, F. (2020). Mitigating effects of temperature variations through probabilistic-based machine learning for vibration-based bridge scour detection. Journal of Civil Structural Health Monitoring, 10(5), 957-972. https://doi.org/10.1007/s13349-020-00427-y
[13] Entezami, A., Shariatmadar, H., & De Michele, C. (2022). Non-parametric empirical machine learning for short-term and long-term structural health monitoring. Structural Health Monitoring, 21(6), 2700-2718. https://doi.org/10.1177/14759217211069842
[14] Bao, Y., Guo, Y., & Li, H. (2020). A machine learning–based approach for adaptive sparse time–frequency analysis used in structural health monitoring. Structural Health Monitoring, 19(6), 1963-1975. https://doi.org/10.1177/1475921720909440
[15] Gunes, B. (2022). One-class machine learning approach for localized damage detection. Journal of Civil Structural Health Monitoring, 12(5), 1115-1131. https://doi.org/10.1007/s13349-022-00599-9
[16] Malekloo, A., Ozer, E., AlHamaydeh, M., & Girolami, M. (2022). Machine learning and structural health monitoring overview with emerging technology and high-dimensional data source highlights. Structural Health Monitoring, 21(4), 1906-1955. https://doi.org/10.1177/14759217211036880
[17] Liu, W., Tang, Z., Lv, F., & Chen, X. (2021). Multi-feature integration and machine learning for guided wave structural health monitoring: Application to switch rail foot. Structural Health Monitoring, 20(4), 2013-2034. https://doi.org/10.1177/1475921721989577
[18] Li, S., Niu, J., & Li, Z. (2021). Novelty detection of cable-stayed bridges based on cable force correlation exploration using spatiotemporal graph convolutional networks. Structural Health Monitoring, 20(4), 2216-2228. https://doi.org/10.1177/1475921720988666
[19] Zhang, L., Qiu, G., & Chen, Z. (2021). Structural health monitoring methods of cables in cable-stayed bridge: A review. Measurement, 168, 108343. https://doi.org/10.1016/j.measurement.2020.108343
[20] Wang, G., Shao, J., Xu, W., Dong, Z., Chen, B., & Dai, X. (2022). Real-time quantitative evaluation on the cable damage of cable-stayed bridges using the correlation between girder deflection and temperature. Structural Health Monitoring, 21(4), 1483-1500. https://doi.org/10.1177/14759217211035048
[21] Ren, J., Zhang, B., Zhu, X., & Li, S. (2022). Damaged cable identification in cable-stayed bridge from bridge deck strain measurements using support vector machine. Advances in Structural Engineering, 25(4), 754-771. https://doi.org/10.1177/13694332211049996
[22] Fan, Z. Y., Huang, Q., Ren, Y., Zhu, Z. Y., & Xu, X. (2020). A cointegration approach for cable anomaly warning based on structural health monitoring data: An application to cable-stayed bridges. Advances in Structural Engineering, 23(13), 2789-2802. https://doi.org/10.1177/1369433220924793
[23] Zhao, W., Zhang, G., & Zhang, J. (2020). Cable force estimation of a long-span cable-stayed bridge with microwave interferometric radar. Computer-Aided Civil and Infrastructure Engineering, 35(12), 1419-1433. https://doi.org/10.1111/mice.12557
[24] Wang, Y., Yang, D. H., & Yi, T. H. (2022). Displacement model error-based method for symmetrical cable-stayed bridge performance warning after eliminating variable load effects. Journal of Civil Structural Health Monitoring, 12(1), 81-99. https://doi.org/10.1007/s13349-021-00529-1
[25] Kim, S., Kim, H. K., & Spencer, B. F. (2022). Automated damping identification of long-span bridge using long-term wireless monitoring data with multiple sensor faults. Journal of Civil Structural Health Monitoring, 12(2), 465-479. https://doi.org/10.1007/s13349-022-00556-6
[26] Jiang, H., Wan, C., Yang, K., Ding, Y., & Xue, S. (2022). Continuous missing data imputation with incomplete dataset by generative adversarial networks–based unsupervised learning for long-term bridge health monitoring. Structural Health Monitoring, 21(3), 1093-1109. https://doi.org/10.1177/14759217211021942
[27] Du, A., & Ghavidel, A. (2022). Parameterized deep reinforcement learning-enabled maintenance decision-support and life-cycle risk assessment for highway bridge portfolios. Structural Safety, 97, 102221. https://doi.org/10.1016/j.strusafe.2022.102221
[28] Marvi, Z., & Kiumarsi, B. (2021). Safe reinforcement learning: A control barrier function optimization approach. International Journal of Robust and Nonlinear Control, 31(6), 1923-1940. https://doi.org/10.1002/rnc.5132

[29] Zhang, Z., & Sun, C. (2021). Structural damage identification via physics-guided machine learning: a methodology integrating pattern recognition with finite element model updating. Structural Health Monitoring, 20(4), 1675-1688. https://doi.org/10.1177/1475921720927488