Anomaly Detection in Structural Health Monitoring with Ensemble Learning and Reinforcement Learning

Nan Huang*
Jiangsu University of Science and Technology, Zhenjiang 212000, Jiangsu, China

***Abstract*—This research introduces a novel approach for improving the analysis of Structural Health Monitoring (SHM) data in civil engineering. SHM data, essential for assessing the integrity of infrastructures like bridges, often contains inaccuracies because of sensor errors, environmental factors, and transmission glitches. These inaccuracies can severely hinder identifying structural patterns, detecting damages, and evaluating overall conditions. Our method combines advanced techniques from machine learning, including dilated convolutional neural networks (CNNs), an enhanced differential equation (DE) model, and reinforcement learning (RL), to effectively identify and filter out these irregularities in SHM data. At the heart of our approach lies the use of CNNs, which extract key features from the SHM data. These features are then processed to classify the data accurately. We address the challenge of imbalanced datasets, common in SHM, through a RL-driven method that treats the training procedure as a sequence of choices, with the network learning to distinguish between less and more common data patterns. To further refine our method, we integrate a novel mutation operator within the DE framework. This operator identifies key clusters in the data, guiding the backpropagation process for more effective learning. Our approach was rigorously tested on a dataset from a large cable-stayed bridge in China, provided by the IPC-SHM community. The results of our experiments highlight the effectiveness of our approach, demonstrating an Accuracy of 0.8601 and an F-measure of 0.8540, outperforming other methods compared in our study. This underscores the potential of our method in enhancing the accuracy and reliability of SHM data analysis in civil infrastructure monitoring.**

Keywords*—Structural health monitoring; Anomaly detection; reinforcement learning; differential equation; imbalanced classification

I. Introduction

SHM is a key method for overseeing civil infrastructure, offering insights into structural loads, performance, responses, and future behavior predictions. SHM's widespread adoption has led to a significant increase in data generation; for example, China's Sutong Bridge, with its 785 sensors, produces 2.5 TB of data annually [1, 2]. Analyzing this vast amount of SHM data is challenging due to various anomalies caused by sensor errors, system failures, environmental factors, and more. These issues, compounded by data from significant events like earthquakes or accidents, can jeopardize the accuracy of structural analysis and the predictive power of SHM systems [3].

Implementing sensor-driven SHM methods generates large amounts of sequential data, complicating manual analysis and anomaly detection. Variations in this data may result from diverse factors such as weather, vehicle overloads, accidents, or unexpected events. It is crucial to recognize that not all anomalies indicate structural issues; some may stem from sensor errors, calibration issues, noise, or transmission problems. To tackle these anomalies, solutions can be applied at both hardware and software levels. While hardware solutions like using wired data channels, extra sensors, or self-validating sensors are effective, they are often costly. Therefore, there is a growing preference for advanced data preprocessing techniques specifically designed for anomaly detection.

SHM faces the challenge of data imbalance, where class instances vary significantly in number. To tackle this issue, two approaches are used: data-centric and algorithm-based. Data-centric strategies, such as under-sampling, over-sampling, and hybrid methods, aim to balance class distribution. Notably, the synthetic minority oversampling technique (SMOTE) [4] creates new minority class instances by linear interpolation, while NearMiss [5] under-samples the majority class using a nearest neighbor algorithm. However, over-sampling can lead to overfitting, and under-sampling may lose critical information. Algorithmic approaches focus on emphasizing the underrepresented class. These include modifying ensemble learning, altering decision thresholds, and employing cost-sensitive learning strategies. Cost-sensitive methods treat classification as cost minimization, assigning higher misclassification costs to minority cases. Ensemble methods combine multiple classifiers for a final decision, and threshold adjustment methods tweak the decision threshold during testing. These techniques aim to effectively balance accuracy and information retention in SHM data classification [6].

Furthermore, the incorporation of deep learning methodologies can serve as an avenue to address the challenge of imbalanced classification [7, 8]. Deep Reinforcement Learning (DRL) emerges as a promising solution to handle imbalanced data due to its distinct attributes. By employing a reward mechanism, DRL can assign augmented importance to the minority class, either by imposing stricter penalties for misclassifying instances from the minority class or by offering greater rewards for accurately identifying them. This approach actively counters the bias that conventional techniques display towards the majority class. The advantages of DRL extend beyond the mere balancing of class distribution. It also enriches the visibility of crucial patterns, particularly those associated with the minority class, by effectively filtering out noisy data. DRL's ability to unearth significant yet often overlooked features within the data contributes to the development of a more robust and efficient model [9].

The initial weight configuration in neural networks is crucial for training in SHM prediction. Traditional training, often using gradient-driven algorithms like backpropagation, typically starts with randomly assigned weights. However, the initial weight selection greatly impacts the training's efficiency and outcome. Careful consideration of the initialization strategy is essential for effective training and accurate SHM prediction. One effective approach is population-based training, where the best solution from a range of generated models is chosen as the starting point for the neural network. This method helps avoid the common issue of getting stuck in local optima, prevalent in standard training methods. Notably, simple evolutionary algorithms have shown effectiveness on par with stochastic gradient descent in neural network training [10].

DE [11] is a popular population-based optimization algorithm widely used in solving optimization problems, particularly effective for weight initialization in machine learning. DE offers several advantages: it ensures a broad exploration of the solution space, preventing entrapment in local optima and leading to better weight configurations. It updates weights iteratively based on the difference between current and target solutions, promoting faster convergence and improved performance. DE is also resilient to noise in fitness assessments, adeptly handling data uncertainties during weight initialization and providing stable initial weight settings. Furthermore, DE's flexibility and adjustability in weight initialization permit tailoring to particular problem areas, like establishing weight limits or integrating previous insights. This versatility improves DE's capability to initiate weights that are aptly matched to the distinct learning challenges being addressed.

This study investigates a novel approach combining a RL-based training algorithm with an advanced DE technique, specifically designed for SHM of bridges. It focuses on detecting anomalies in time-series sensor data from a major cable-stayed bridge in China. The data is divided into seven categories: normal, trend, square, missing, minor, drift, and outlier, with 'normal' being the most frequent. To overcome the issue of class imbalance in the dataset, the research introduces a framework that treats classification as a series of strategic decisions. In each iteration, an agent assesses a training sample (environmental state) and makes classification decisions, earning rewards or penalties based on the outcome. Classes with fewer samples are assigned higher rewards, encouraging accurate identification of less common anomalies. Additionally, the study integrates a unique mutation operator based on clustering principles within the DE framework to improve the backpropagation (BP) process. This operator identifies dominant clusters in the DE population and implements a novel approach for creating potential solutions. The key contributions of this research lie in its innovative approach to class imbalance, decision-making process in classification, and enhanced training methodology through integrating RL, DE, and BP processes:

1) We present an innovative RL-based method specifically designed to address the inherent challenges of imbalanced classification in SHM.

2) The approach integrates a unique reward system that reinforces accurate decisions while penalizing incorrect ones. By allocating enhanced rewards to the less represented class, we directly address the challenge of data skewness, encouraging the algorithm to appropriately focus on lesser-known data. This strategic maneuver contributes to a more fair and balanced classification procedure.

3) To extract deeper insights from images and refine the classification decision-making process, we employ a fusion of CNN models. This approach enhances the representation of features, resulting in improved accuracy and robustness in classification efforts.

4) We have developed an enhanced DE algorithm to initialize weights in the proposed model efficiently. This tactic aids in identifying a promising region for initiating the BP algorithm within the model.

The structure of this document is as follows: Section II details a review of relevant literature, and Section III provides an overview of the key dataset utilized in this study. Section IV delves into the proposed strategy, elucidating the core methodology in depth. Section V unfolds the empirical outcomes and their subsequent dissection. Concluding observations and potential avenues for future inquiry are encapsulated in Section VI.

II. Literature Review

Artificial Intelligence (AI) techniques bring forth the capability to uncover patterns within time-series data with no prior comprehension of the underlying structural architecture. These methodologies involve exploring either the time or frequency domain of the data, extracting pertinent characteristics through statistical evaluations, or employing signal processing tools like the Fourier and wavelet transforms, as well as the Hilbert-Huang and Shapelet transforms. On the other hand, deep learning (DL) algorithms possess the ability to autonomously extract significant attributes by interpreting time-frequency data as visual inputs within a CNN framework. However, it is important to acknowledge that DL-centric approaches, while potent, demand substantial computational resources and cause meticulous fine-tuning of hyperparameters [12].

In order to tackle irregularities within SHM data, Pan et al. [13] presented an approach rooted in transfer learning. They employed a deep neural network to discern and rectify aberrant data, enhancing the accuracy of bridge evaluations. Samudra et al. [12] devised a comprehensible framework rooted in decision trees, employing random forest classifiers to categorize acceleration data in the realm of SHM. This approach, boasting a remarkable 98% accuracy, emerges as an economically viable avenue for gauging infrastructure state. Li et al. [14] outlined a strategy to elevate the efficacy of anomaly detection within bridge SHM systems. Employing strategies like data augmentation, feature dimension reduction, and a two-stage deep convolutional neural network, they achieved an elevated level of recognition accuracy. Tang et al. [3] presented an innovative anomaly detection method catering to SHM, employing a CNN that transforms time series data into visual representations. This approach achieves precise identification of diverse pattern anomalies, scaling effectively and bolstering accuracy. Ye et al. [15] proposed a technique rooted in deep learning for identifying data anomalies within SHM systems. By deploying time-frequency analysis and CNNs, they translated SHM data into RGB images, subsequently classified through a Google network. Green et al. [16] introduced a new way to use Bayesian techniques in the analysis of inclinometer data for SHM. This method allows for the detection of anomalies, forecasting, and quantifying uncertainties, leading to better risk assessment and cost reduction.

Moreover, the framework has the potential to be applied in various engineering fields beyond inclinometers. Boccagna et al. [17] suggested an AI approach for monitoring structural health in almost real-time, using unsupervised deep learning. By preprocessing data and utilizing artificial neural network autoencoders, the technique effectively identifies anomalies, surpassing current methods and demonstrating encouraging outcomes. Lei et al. [18] proposed a residual attention network (RAN) to detect abnormal data in measured structures. The RAN incorporates attention mechanisms and residual learning to enhance classification accuracy and efficiency. It achieved exceptional performance and generalization on datasets from an arch bridge and a cable-stayed bridge, surpassing existing models in terms of multi-classification and accuracy. Yang and Nagarajaiah [19] introduced a principled, independent component analysis approach to reduce faulty data during data transmission; then, they achieved reliable data transfer and image restoration using compression sensing methods [20]. Yang and Nagarajaiah [21] employed the principal component pursuit method to detect and minimize burst noise in ambient vibration response. They also introduced a data management and processing framework based on sparsity rank and low-rank techniques. Park et al. [22] utilized transmission errors and ensemble empirical mode decomposition to identify anomalies such as gear teeth spalls and cracks in rotating machinery.

III. Dataset Description

In this study, our primary focus centers on the detection of specific deviations - including trends, squares, omissions, minor variances, drifts, and outlier - within the acceleration time series derived from a lengthy cable-stayed bridge in China. The IPC-SHM community [23] provides access to curated data from this bridge. An overview of these anomalies, distilled from the bridge's measured data, is concisely summarized in Table I.

Besides the irregularities mentioned, it is vital to acknowledge that acceleration sensors, particularly those affixed to structures with potential vulnerabilities, adeptly detect a broad range of atypical patterns. These include offsets, characterized by sudden, noticeable jumps in response, and gains, marked by a slow, consistent increase in response over time. Furthermore, the sensors can identify precision deterioration, where the response shows erratic fluctuations, and complete failures, which result in a response akin to the randomness of white noise in the frequency domain. Recognizing and interpreting these additional types of deviations are crucial for comprehensive structural health monitoring, as they can provide early warning signs of more significant issues or impending failures.

The dataset under study contains a thorough record of acceleration data over a month, meticulously gathered from 38 strategically placed accelerometers across the bridge. For detailed analysis, this data is segmented into individual hourly time series, leading to an extensive collection of 28,272 such series. This figure is calculated considering the number of sensors, the days in the month, and the daily time cycle. Given that the accelerometers recorded at a rate of 20 Hz, the total data volume reaches an astonishing 2×10^9 data points, a multiplication of the number of time series, seconds in an hour, and the sampling rate. This vast dataset offers a rich source for in-depth analysis, allowing for the examination of minute changes and patterns over time, providing a comprehensive understanding of the bridge's dynamic behavior under various conditions.

The subsequent classification task systematically organizes these 28,272 time series responses into seven distinct categories. This includes the ‘normal’ set and six other types of anomalies: trend, square, missing, minor, drift, and outlier. A detailed chart in Table I itemizes the precise distribution of time series across each category within this dataset. This categorization is crucial for identifying the predominant anomalies and understanding their relative occurrences. It aids in developing targeted strategies for monitoring and maintenance, ensuring focused attention on the most critical or frequently occurring issues, enhancing the overall efficiency and effectiveness of the structural health monitoring process.

TABLE I. Description of the Anomalies

Class	Description	Count
Normal	The time-domain response showcases balance, while multiple resonance peaks can be observed in the frequency-domain response.	13575
Trend	A discernible trajectory is apparent in the time-domain response, and a distinctive peak value is identified in the frequency-domain response.	5778
Square	The time-domain response mirrors a square wave.	2996
Missing	The bulk of the time-domain response is absent.	2942
Minor	Compared to standard category data, the time-domain response exhibits a notably reduced magnitude.	1775
Drift	The time-domain response varies unpredictably, either exhibiting arbitrary shifts or increasingly diverging over time.	679
Outlier	The time-domain response contains singular or multiple pronounced protrusions.	527

IV. Model Architecture

Fig. 1 illustrates the intricacies of our innovatively developed model, meticulously engineered to boost the efficiency of anomaly detection. This model is specifically tailored to tackle issues like uneven distribution of classes and the critical importance of accurately setting initial weights. By integrating the DE algorithm with PPO, our model adeptly circumvents common hurdles encountered in standard models.

Conventional methods often miss a systematic strategy for establishing initial weights, resulting in reduced learning speeds and a propensity to settle on less-than-ideal outcomes. Our technique leverages DE to provide a diverse spectrum of initial weights. This variety helps the model to bypass smaller optima and more efficiently converge on more comprehensive solutions. Expanding on this, the utilization of DE in our method is not just about broadening the range of initial weights, but also about introducing a dynamic and adaptive way of initializing these weights. This adaptability is crucial in complex models where the landscape of solutions is vast and varied. By starting from a more helpful position in the solution space, our method enhances the model's ability to navigate through this landscape, leading to quicker and more effective convergence. Furthermore, this approach also contributes to the robustness of the model, making it less susceptible to the challenges posed by different data distributions and complexities inherent in various learning tasks. Our method does not just improve the efficiency of the learning process, but also broadens its applicability and effectiveness across a range of scenarios.

Additionally, the RL element in our framework is thoughtfully structured to significantly favor the accurate identification of the minority class, underscoring these crucial predictions. This represents a significant advancement over conventional supervised learning approaches, which often struggle with insufficient data representation across various classes. The dynamic nature of policy learning in RL encourages a fairer approach to decision-making, leading to enhanced strategies for identifying underrepresented categories. The flexibility of RL in our setup sets it apart from orthodox methodologies, equipping it with the essential capabilities to efficiently address the typical challenges found in conventional classification techniques employed in anomaly detection.

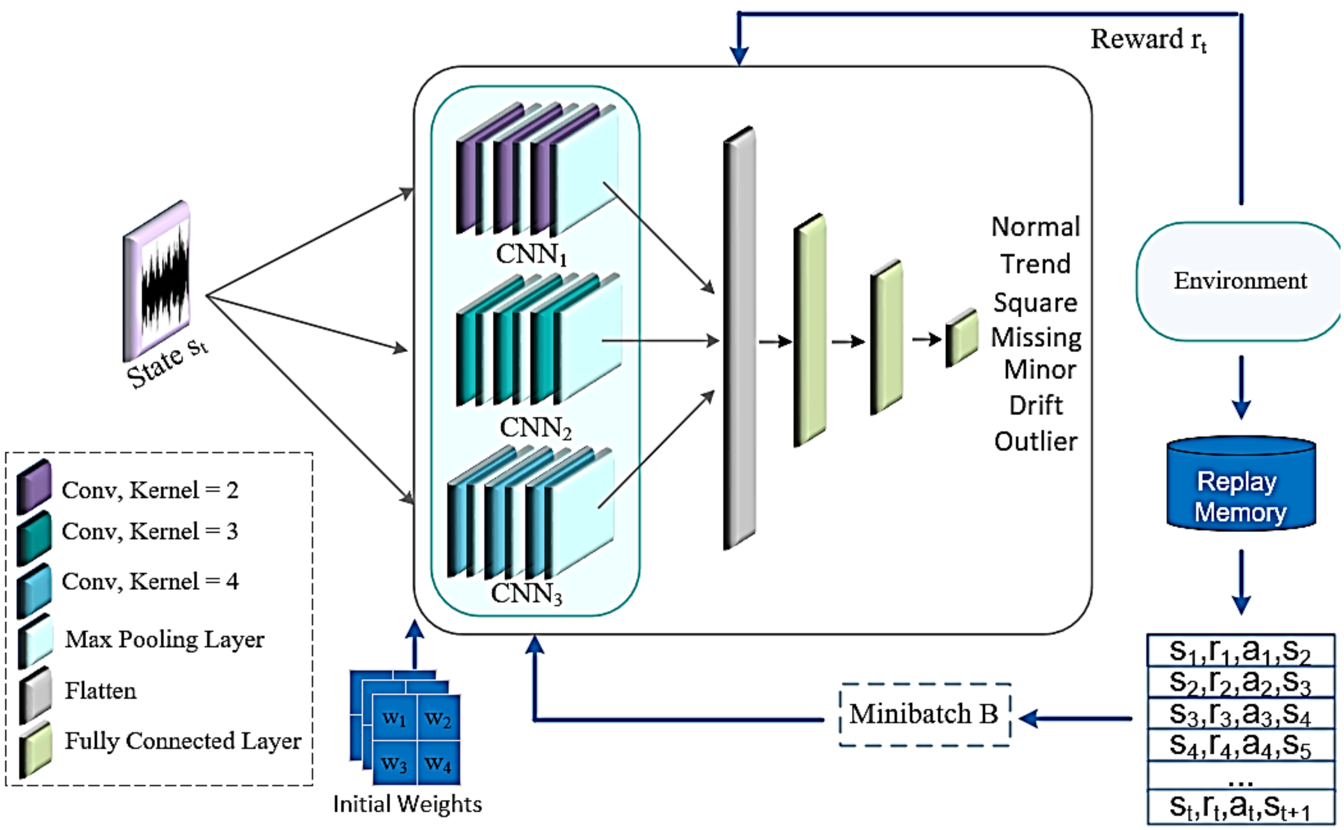

Fig. 1. Our model pipeline encompasses a sequence of procedures.

A. Pre-training Phase

Deep models depend heavily on the initialization of deep network weights. If the initial values are not accurate, it can lead to convergence issues in the model. The first stage in this paper is to set the CNN and feed-forward neural network weights. We offer a more effective DE approach, which incorporates the power of a clustering technique and an innovative fitness function to optimize its performance. In our improved DE algorithm, we utilize a mutation and updating scheme based on clustering to enhance the optimization performance. Complex architectures rely significantly on the initial setup of deep network parameters. Inaccurate starting points may cause the model to struggle with convergence. The initial step in our study involves configuring the weights for the CNN and the feed-forward neural networks. We propose an enhanced DE method that integrates the effectiveness of a clustering algorithm and a novel fitness function to boost its efficiency. In our refined DE technique, we employ a mutation and refresh strategy centered on clustering to improve the optimization process. Extending this, our approach takes into account the intricacies of deep learning architectures, ensuring that the weight initialization is not only randomly, but strategically influenced by the underlying data structure. The clustering-based mutation strategy allows for a more targeted and data-driven adjustment of weights, which is particularly useful in navigating the high-dimensional spaces typical in deep learning. This method aids in avoiding local minima and accelerates the convergence of the model.

The mutation mechanism, influenced by studies referenced in [24], identifies a promising region in the exploration domain. Employing the k-means clustering technique, the existing group P is segregated into k segments, each representing a unique portion of the exploration zone. A random integer chosen from the interval $[2, \sqrt{N}]$ dictates the cluster count. The cluster deemed most optimal possesses the lowest mean fitness value across its gathered samples post-clustering. Expanding on this, our method enhances the search strategy within the algorithm. By dividing the population into clusters, we can pinpoint specific regions in the search space that hold potential for better solutions. This clustering not only focuses the search but also adds a layer of precision in identifying promising areas, thereby increasing the efficiency of the mutation process. Additionally, the number of clusters is dynamically determined based on the population size, allowing for a flexible and adaptable approach to clustering. This adaptability is crucial in dealing with diverse problems and varying sizes of search spaces. The concept of assessing the perfection of a cluster based on its average fitness introduces a competitive element among the clusters, driving the algorithm to favor areas of the search space that show higher potential for optimal solutions. Furthermore, our approach refines the selection process within each cluster. After identifying the most optimal cluster, we focus on fine-tuning the solutions within this cluster, leveraging the collective intelligence of the group. This targeted mutation within the most promising cluster ensures the algorithm does not just wander aimlessly across the entire search space but makes informed, strategic moves towards areas more likely to yield superior results.

The clustering-based approach outlines the proposed mutation:

$$\overrightarrow{v_i^{clu}} = \overrightarrow{win_g} + F(\overrightarrow{x_{r_1}} - \overrightarrow{x_{r_2}}) \tag{1}$$

where, $\overrightarrow{x_{r_1}}$ and $\overrightarrow{x_{r_2}}$ represent two candidate solutions randomly selected from the current population while $\overrightarrow{win_g}$ corresponds to the best solution within the promising region. It is important to note that $\overrightarrow{win_g}$ may not be the best solution for the entire population.

Following the creation of M new solutions via mutation grounded in clustering, the current population undergoes an update under GPBA [25]. The GPBA is an innovative optimization approach that meticulously navigates through the search space, assessing the effectiveness of solutions against a series of established patterns. In practical application, GPBA changes the population by meticulously choosing and substituting individuals according to their performance indicators. By employing these patterns as navigational aids, GPBA refines its quest for the best solutions, often resulting in a more streamlined and efficient path to the global optimum. This method's real strength lies in its structured approach to the exploration of the search space. By using gradient patterns, the algorithm can intelligently predict the direction in which improvements can be made, rather than relying on random or exhaustive search methods. This predictive capability is helpful in complex optimization scenarios, characterized by vast search spaces and elusive optimal solutions. It allows the algorithm to bypass fewer promising regions of the search space, focusing its efforts on areas more likely to yield fruitful results. Furthermore, the GPBA's adaptability to different optimization problems adds to its versatility. Whether the task involves continuous or discrete variables, linear or non-linear relationships, the GPBA can be tailored to suit the specific characteristics of the problem. This adaptability is achieved through the customization of its pattern-based search mechanisms, which can suit various problem structures and complexities. Besides its efficiency in finding solutions, the GPBA also offers improved computational speed compared to more traditional optimization methods. This is beneficial in real-time applications or scenarios where time is a critical factor. The algorithm's ability to quickly converge to an optimal solution without sacrificing accuracy makes it an attractive choice for a wide range of optimization tasks.

The process unfolds as follows:

- Selection: To initiate the algorithm, generate k random individuals that will function as the initial points.
- Generation: Produce a set of M solutions using mutation based on clustering and denote it as v^{clu}.
- Replacement: Choose M solutions randomly from the current population to form set B.
- Update: Select the top M solutions from the combined groups v^{clu} and B to create a new group B'. The refreshed population is derived by merging members of set P not included in B with those from set B' ($(P - B) \cup B'$).

B. DRL

DRL stands as a formidable approach in the domain of deep learning. Within this framework, an intelligent agent engages dynamically with its environment, aiming to maximize its cumulative rewards. This flexible and adaptive learning mechanism empowers the agent to make a series of decisions, often in the face of uncertainty, which has profound applications across a wide spectrum of domains, including but not limited to robotics, healthcare, and finance [26]. The prowess of DRL becomes evident in tasks that require sequential decision-making and the ability to adapt to unforeseen and evolving circumstances. Its capacity to handle complex activities that unfold over time, adjusting its strategies and responses as needed, underscores its versatility and broad applicability in addressing real-world challenges. DRL's ability to learn from interactions with the environment, optimize decision-making processes, and navigate through dynamic scenarios positions it as a valuable tool for a wide range of applications, making it a compelling area of research and development in artificial intelligence.

In categorization-related tasks, a major challenge lies in handling datasets with imbalanced distributions, where one category is markedly more dominant than others. This disproportion might cause skewed educational outcomes, as standard classification approaches often lean towards the predominant group, leading to subpar identification of the less represented categories. Under such conditions, DRL stands out as a superior strategy for educating neural networks over conventional approaches. DRL addresses the problem of lopsided categorization by employing a system based on rewards [27]. Through carefully allocating incentives, it shifts the agent's attention towards instances belonging to the underrepresented categories, thus improving the detection of these rarer classes. The reward-centric model of DRL promotes a comprehensive decision-making process, prioritizing the discovery and classification of rare events or infrequently occurring categories.

In the realm of DRL, the primary goal of the agent is to select actions that optimize prospective benefits. The aggregation of rewards for forthcoming situations, symbolized by the reward value, gradually decreases over time, influenced by the discount rate γ, as illustrated in Eq. (2). In this formula, T corresponds to the concluding time-step of an episode [28].

$$R_t = \sum_{t'=t}^{T} \gamma^{t'-t} r_{t'} \tag{2}$$

where, R_t represents the cumulative reward starting from time t, and $r_{t'}$ denotes the reward received at time t'. Q-values, representing the quality of state-action interactions, denote the anticipated outcome of policy π upon executing action a within state s. This is computed as depicted in Eq. (3).

$$Q^{\pi}(s,a) = E[R_t | s_t = s, a_t = a, \pi] \tag{3}$$

The most optimal action-value function, represented as the highest anticipated reward among all approaches after witnessing state s and performing action a, is calculated as depicted in Eq. (4).

$$Q^{*}(s,a) = max_{\pi}\, E[R_t | s_t = s, a_t = a, \pi] \tag{4}$$

The function carries out the Bellman equation [29], which states that the supreme anticipated outcome for a particular maneuver is the sum of the benefits from the present maneuver and the utmost anticipated outcome from forthcoming maneuvers in the next instance. This concept is exemplified in Eq. (5).

$$Q^{*}(s,a) = E[r + \gamma\, max_{a'} Q^{*}(s',a') | s_t = s, a_t = a] \tag{5}$$

The computation of the ideal action-value function is methodically executed using the Bellman equation, as illustrated in Eq. (6).

$$Q_{i+1}(s,a) = E[r + \gamma\, max_{a'} Q_i(s',a') | s_t = s, a_t = a] \tag{6}$$

During the learning stage, as the network experiences state s, it generates a state-specific action. Subsequently, the system provides a reward r and transitions to the next state s'. These components are combined into a set (s,a,r,s'), subsequently stored in memory M. Groups of such sets, termed Batches B, are selected for performing gradient descent. The method for calculating loss is detailed in Eq. (7).

$$L_i(\theta_i) = \sum_{(s,a,r,s') \in B} (y - Q(s,a;\theta_i))^2 \tag{7}$$

Here, θ symbolizes the model's weights, while y indicates the approximated objective for the Q function, evaluated as the summation of the reward linked with the state-action pair and the reduced maximum Q value in future instances, as illustrated in Eq. (8).

$$y = r + \gamma\, max_{a'} Q(s',a';\, \theta_{k-1}) \tag{8}$$

It is important to recognize that the Q value assigned to the terminal state is initialized at zero. The gradient's magnitude for the loss function during the i-th iteration is ascertainable through Eq. (9).

$$\nabla_{\theta_i} L(\theta_i) = -2 \sum_{(s,a,r,s') \in B} (y - Q(s,a;\theta_i)) \nabla_{\theta_i} Q(s,a;\theta_i) \tag{9}$$

Through the execution of a gradient descent iteration on the loss function, adjustments are made to the model's weights under Eq. (10). This modification endeavors to lessen the discrepancy, where α denotes the learning rate dictating the extent of advancement within the optimization procedure.

$$\theta_{i+1} = \theta_i + \alpha \nabla_{\theta_i} Q(s,a;\theta_i) \tag{10}$$

1) Problem formulation: Within this paper, the application of the RL algorithm is directed towards the field of SHM. The ensuing explanation delineates the method's functioning and interpreting each component:

- State $\boldsymbol{s_t}$: This matches the image captured at the temporal interval t. Here, an image refers to a graphical representation of the time series data. This image is composed of plots or graphs that visually depict the various anomalies identified in the acceleration time series of the bridge.
- Action $\boldsymbol{a_t}$: The categorization executed on the image is regarded as an action. This signifies a choice carried out by the network, grounded in its prevailing comprehension of the objective.
- Action $\boldsymbol{r_t}$: A reward is furnished for every categorization, designed to steer the network towards accurate categorization. The formulation of this remuneration process is expressed as:

$$r_t(s_t, a_t, y_t) = \begin{cases} +1\,, a_t = y_t \text{ and } s_t \in D_O \\ -1\,, a_t \neq y_t \text{ and } s_t \in D_O \\ \lambda\,, a_t = y_t \text{ and } s_t \in D_N \\ -\lambda\,, a_t \neq y_t \text{ and } s_t \in D_N \end{cases} \tag{11}$$

In this context, $D_N = \{Noraml\}$, and $D_O = \{\text{Trend, Square, Missing, Minor, Drift, Outlier}\}$ indicates the minority classes. Accurate or erroneous classification of a case from the prevalent category leads to an incentive or penalty of +λ or -λ, respectively. This outlined method compels the network to focus on accurately identifying instances from the less frequent class by allocating a higher absolute value to the reward. Concurrently, the incorporation of the normal class and the flexible reward parameter within the range of 0<λ<1 adds complexity to the reward scheme. This allows for refined adjustment of the network's focus between the more and less prevalent classes.

V. Empirical Evaluation

In the meticulous assessment stage, a detailed and exhaustive analysis was carried out, contrasting our suggested model with six distinct deep-learning contenders, namely TransAnoNet [13], AnoSegNet [14], WaveletCNN [15], GAN-VAE [30], CNN-MIAD [31], and VibroCNN. This evaluation aimed to provide an all-encompassing insight into the strengths of our model vis-à-vis established methods. Moreover, we delved into different versions of our model by introducing three alternative variants for evaluation. The initial variant, termed "Proposed without dilated convolution," was based on a similar foundational architecture to our original model, yet did not incorporate dilated convolution. The subsequent variant, designated as "Proposed without RL," excluded the reinforcement learning component from the classification procedure. The third altered version, named "Proposed without DE," employed random initialization for the weights. We appraised these models using standard performance indicators, focusing specifically on measures like the F-measure and the geometric mean due to their proven effectiveness in tackling imbalanced datasets. The findings, detailed in Table II, resoundingly affirm the preeminence of our proposed model over all competing models, including those previously recognized as industry standards like AnoSegNet and TransAnoNet. Across every evaluative criterion, our model demonstrated consistent superiority over its rivals. Noteworthy accomplishments involve a marked error reduction exceeding 9% in the F-measure and surpassing 8% in the G-means indices. These notable advancements highlight the efficacy of our model in surmounting the difficulties associated with imbalanced datasets and its adeptness at furnishing more accurate forecasts. In juxtaposing our model with the modified iterations "Proposed without dilated convolution," "Proposed without RL," and "Proposed without DE," the indispensability of incorporating dilated convolution, RL, and DE methods becomes clear. Our model manifested a notable error rate reduction, approximately 5.35%, in comparison to its counterparts. This outcome accentuates the critical impact that the amalgamation of dilated convolution, RL, and DE strategies has in boosting the model's performance, thus solidifying their role as catalysts in the evolution of deep learning models.

In Fig. 2, we present the receiver operating characteristic (ROC) curves corresponding to the methodologies outlined in Table II. The area under the curve (AUC) serves as a pivotal metric for quantifying the performance of classifiers. An AUC score of 1 signifies impeccable discrimination ability, while a score of 0.5 suggests performance no better than random guessing.

It is worth highlighting that our proposed model emerges as the leader in this analysis, boasting a notable AUC value of 0.60. This solid outcome highlights its remarkable proficiency in accurately differentiating between favorable and unfavorable results, reinforcing the credibility of our approach as a potent predictive tool. Additionally, the "Proposed without RL" approach also demonstrates strong performance, achieving an AUC of 0.57, further affirming its ability to discern between positive and negative instances. In contrast, WaveletCNN and TransAnoNet, which achieve AUC scores of 0.46 and 0.49, respectively, offer less impressive performance. VibroCNN, GAN-VAE, and CNN-MIAD display even less favorable outcomes, with AUC values ranging from 0.43 to 0.45. Particularly, VibroCNN's meager AUC of 0.43, only slightly surpassing random chance, highlights its underwhelming performance. The ROC analysis vividly illustrates the varying degrees of performance among the evaluated methodologies. The exceptional predictive prowess demonstrated by our proposed method, whether in its standalone form or when coupled with RL, underscores the potency of our approach. Furthermore, it establishes a robust foundation for future enhancements and promising applications in the realm of predictive modeling, charting a path toward even more effective methodologies in the prediction domain. This remarkable performance positions our model as a key player in the field of predictive analytics.

TABLE II. Efficiency Indicators of the Suggested System Compared to Rival Advanced Networks for SHM

	Accuracy	**F-measure**	**G-means**
TransAnoNet	0.8104±0.0156	0.7802±0.1053	0.8102±0.0203
AnoSegNet	0.8001±0.2156	0.7371±0.0268	0.7902±0.1236
WaveletCNN	0.8005±0.2130	0.7202±0.0268	0.7906±0.2156
GAN-VAE	0.6801±0.0526	0.5402±0.1236	0.6405±0.0256
CNN-MIAD	0.7703±0.1563	0.6602±0.1036	0.7403±0.1265
VibroCNN	0.6704±0.0501	0.5501±0.0652	0.6462±0.0052
Proposed without dilated convolution	0.7904 ± 0.0517	0.7615 ± 0.1623	0.8014 ± 0.3622
Proposed without RL	0.8315 ± 0.0243	0.8200 ± 0.0417	0.8315 ± 0.0621
Proposed without DE	0.8425 ± 0.0123	0.8345 ± 0.0120	0.8436 ± 0.0505
Proposed	0.8601 ± 0.0384	0.8540 ± 0.0297	0.8760 ± 0.0123

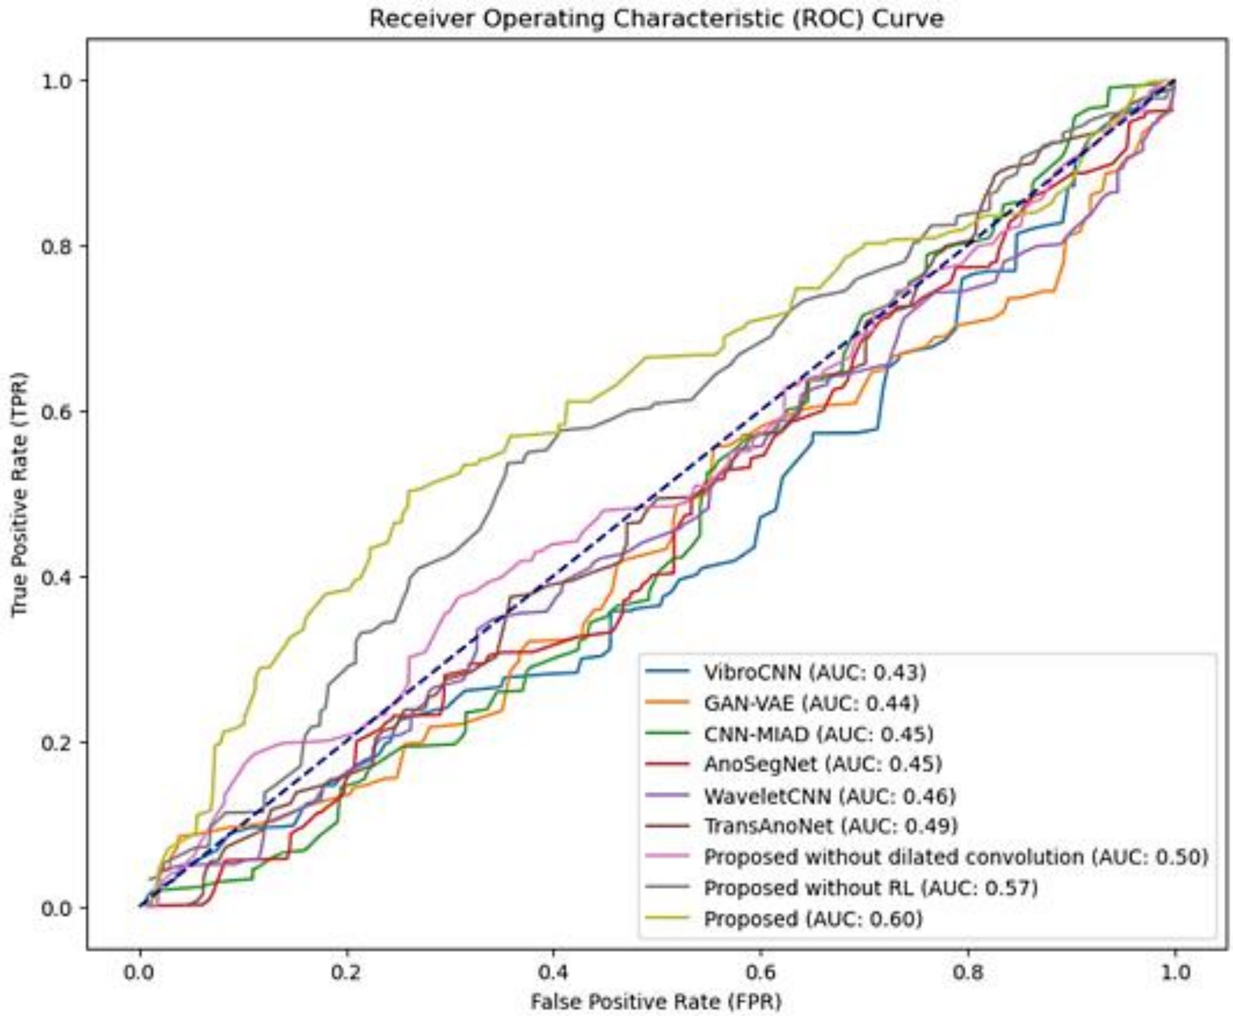

Fig. 2. AUC chart for the suggested approach and alternative comparative techniques.

Fig. 3 showcases the confusion matrices for the proposed model, providing a detailed representation of its classification performance across different categories. From the matrix, we can observe the number of correct predictions (true positives) along the diagonal for each class, which are as follows: 13,271 for 'Normal', 5,363 for 'Trend', 2,608 for 'Square', 2,750 for 'Missing', 1,571 for 'Minor', 539 for 'Drift', and 409 for 'Outlier'. These figures suggest the model is most proficient at identifying the 'Normal' class and least proficient at identifying 'Outlier' instances, which could be because of their lower occurrence in the dataset. The off-diagonal numbers represent the instances where the model misclassified the inputs. For example, there are 76 instances where 'Normal' was incorrectly classified as 'Missing', and 80 instances where 'Square' was mislabeled as 'Normal'. Such misclassifications can diagnose and improve the model's performance, possibly by providing it with more representative training data or refining its feature detection capabilities.

Fig. 4 illustrates the evolution of error dynamics within the proposed model across 500 epochs. Commencing at an initial value of 12, the error undergoes a consistent descent as epochs unfold. This sustained decline signifies the model's progressive learning and enhancement of its predictive capabilities over time. It is significant to observe that the most pronounced decrease in error happens during the early training stages, slowly leveling off as the number of epochs increases. This trend indicates that with ongoing training, the rate of error reduction lessens, signifying a point where further error minimization from extended training becomes less impactful. Near the 425th epoch, a clear steadying of the error rate is observed, consistently hovering around a value of approximately 4.2962 in subsequent epochs. This leveling off of error rates suggests that continued training beyond this juncture is unlikely to result in notable enhancements in the model's forecasting accuracy. This stage may signal that the model has attained a state of convergence, reaching an accuracy level where additional fine-tuning might not bring considerable improvements. Alternatively, this stabilization could also suggest the emergence of overfitting concerns, especially if the model's performance on validation or test datasets ceases to improve. This insight into the error dynamics across epochs not only showcases the model's learning journey but also provides valuable guidance for fine-tuning training duration and preventing potential overfitting scenarios.

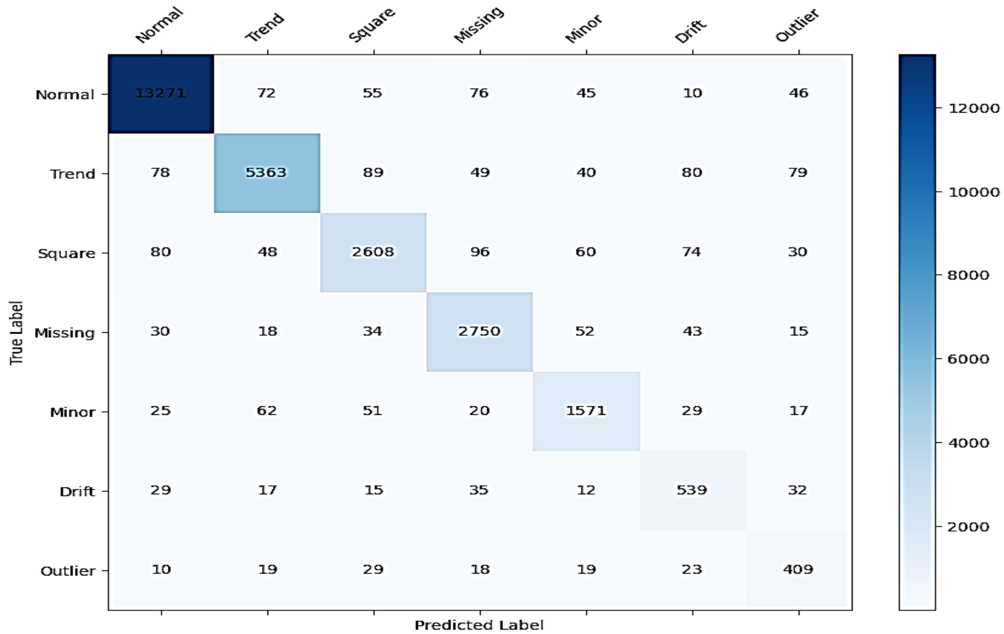

Fig. 3. Comparative confusion matrices for the proposed model.

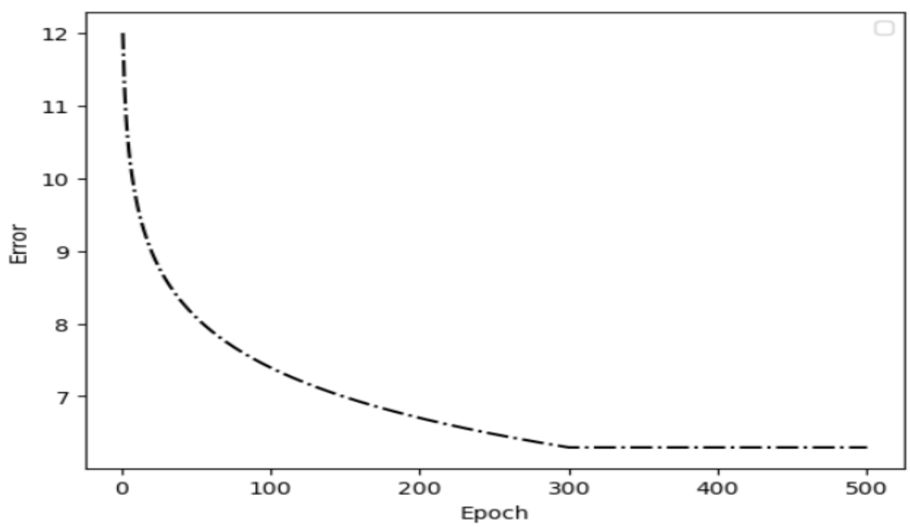

Fig. 4. Comparative diagram of error dynamics.

A. Impact of the Reward Function

The allocation of rewards to both the more common and less frequent categories for accurate and inaccurate classifications is denoted by +1 and ±λ. The particular magnitude of λ is determined by the ratio of frequent occurrences relative to rare events. As this ratio rises, it is expected that the ideal magnitude of λ will diminish proportionally. In order to thoroughly investigate the influence of λ, we executed an extensive assessment of the suggested structure employing diverse λ magnitudes, varying from 0 to 1 in steps of 0.1. Concurrently, the incentive for the more frequent category stayed unchanged. The detailed results are illustrated in Fig. 5. When adjusting λ to 0, the effect of the dominant group turns negligible.

Conversely, with a value of λ = 1, both the more common and less common groups carry equivalent weight. The insights extracted from the analysis indicate that the framework achieves its peak effectiveness when λ is established at 0.7, as observed across all assessed performance indicators. This observation suggests that the ideal λ magnitude lies within the range of zero to one. It's important to recognize that although modulating λ to reduce the impact of the dominant group is essential, configuring it too low might adversely affect the overall effectiveness of the entire structure. The evidence clearly indicates that the choice of λ markedly affects the success of the structure. The suitable λ magnitude depends on the comparative occurrences of more frequent and infrequent events, highlighting the need for careful determination to achieve the best results. This study underscores the intricate interplay between λ and the framework's success, advocating for a balanced choice of λ to strike a harmonious equilibrium between the two categories and foster effective results.

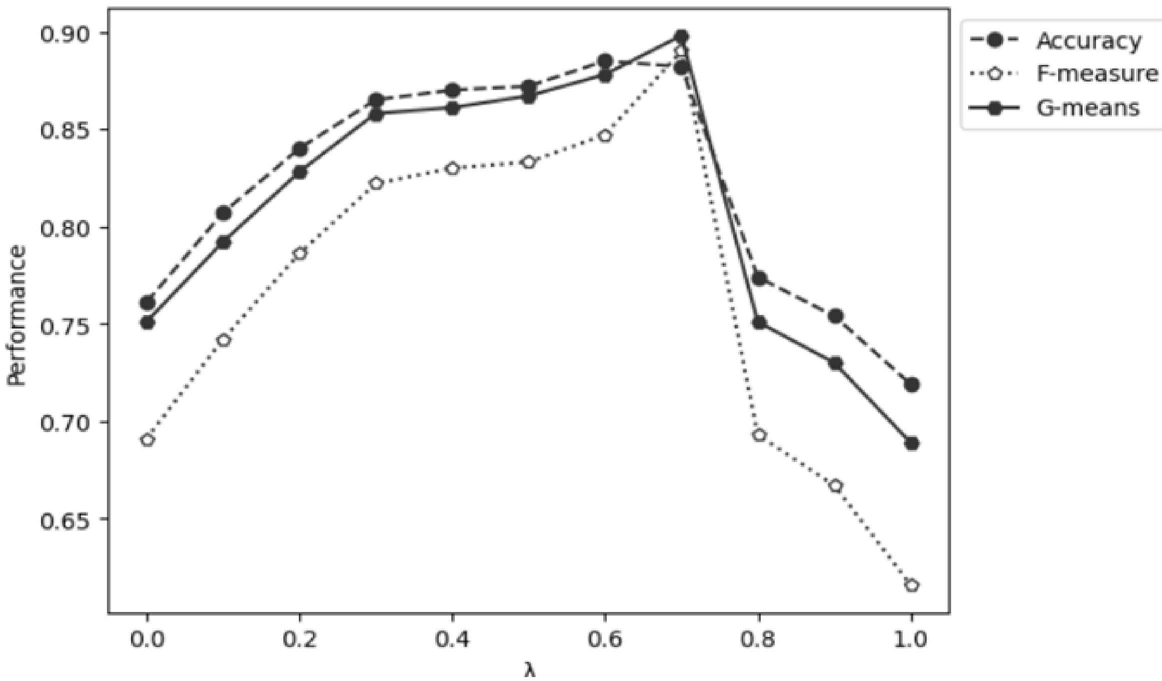

Fig. 5. Evaluation of the performance metrics of the proposed system under various settings of the parameter λ.

B. *Effect of Loss Function*

The landscape of strategies available to combat the complexities arising from data imbalances in machine learning models is vast and diverse. It spans an array of techniques, ranging from the fine-tuning of data augmentation methods to the meticulous selection of aptly suited loss functions. The deliberate choice of an appropriate loss function plays a central role in ensuring the model's capacity to glean valuable insights from the underrepresented class embedded within the dataset. In our quest to unravel the nuances of the varying impacts of distinct loss functions, we embarked on a comprehensive exploration of five distinct contenders: WCE [32], BCE [33], DL [34], TL [35], and CL [36].

Among these contenders, both BCE and WCE have established themselves as widely adopted loss functions, treating positive and negative samples with equal significance. However, it's imperative to recognize that these functions might not be optimally configured to cater to datasets characterized by pronounced imbalances that accentuate the minority class. In stark contrast, DL and TL exhibit superior performance when confronted with skewed datasets, delivering more favorable outcomes for the underrepresented class. Notably, CL emerges as a standout loss function, showcasing its prowess in scenarios where imbalanced data prevails. By skillfully adjusting the weights of the loss function, CL demonstrates its ability to prioritize intricate samples over simpler ones, thereby enhancing its adaptability in the face of challenging data distributions.

Our rigorous experimentation and analysis of these diverse loss functions are presented in meticulous detail in Table III. The outcomes unequivocally affirm the supremacy of CL over TL, leading to a substantial 3.72% reduction in the error rate concerning accuracy and an impressive 3.58% decrease in the F-measure. Nevertheless, it is crucial to underscore that, when benchmarked against the performance of our proposed model, CL exhibits a 1.5% deficit. These findings underscore the paramount significance of making a judicious selection of an appropriate loss function when navigating the intricacies of imbalanced data. Furthermore, they shine a spotlight on the commendable performance of our model in effectively addressing this prevalent and challenging issue in machine learning.

TABLE III. PERFORMANCE EVALUATION METRICS OF THE PROPOSED MODEL AGAINST VARIOUS LOSS FUNCTIONS IN SHM

	Accuracy	F-measure	G-means
WCE	0.7303± 0.0269	0.7105± 0.1204	0.7412± 0.1120
BCE	0.7963± 0.0626	0.7536± 0.1203	0.7963± 0.0103
DL	0.7923± 0.0365	0.7821± 0.0056	0.8103± 0.1123
TL	0.8236± 0.2126	0.8023± 0.0145	0.8352± 0.1035
CL	0.8563± 0.0035	0.8325± 0.0032	0.8523± 0.0039

C. Effect of CNNs

The architecture encompasses an array of CNNs that concurrently derive feature vectors from images. The quantity of CNNs utilized for feature extraction greatly influences the model's efficiency. Inadequate number of CNNs results in inadequate feature extraction, whereas too much may lead to problems such as overfitting or superfluousness. Both scenarios can diminish the model's overall utility. Therefore, carefully selecting the ideal number of CNN feature extractors is crucial. To identify this optimal number, we conducted a thorough and systematic analysis, evaluating the model's performance across a range of 1 to 7 CNN feature extractors. Our aim was to pinpoint the point where the model achieves peak functionality while maintaining a delicate balance between thorough feature extraction and operational efficiency. Our comprehensive experiments, as illustrated in Fig. 6, unequivocally demonstrate that three CNN feature extractors yield the model's best performance. Interestingly, as the number of CNNs increases, the model's performance declines, with six or seven extractors being less effective than using just one. This observed pattern highlights the existence of an optimal number of CNN feature extractors, maximizing the model's ability to capture relevant and discriminative features, ultimately leading to enhanced overall performance. Selecting three CNN feature extractors strikes a harmonious balance, optimizing the model's ability to extract essential information from input imagery, boosting its efficiency and effectiveness.

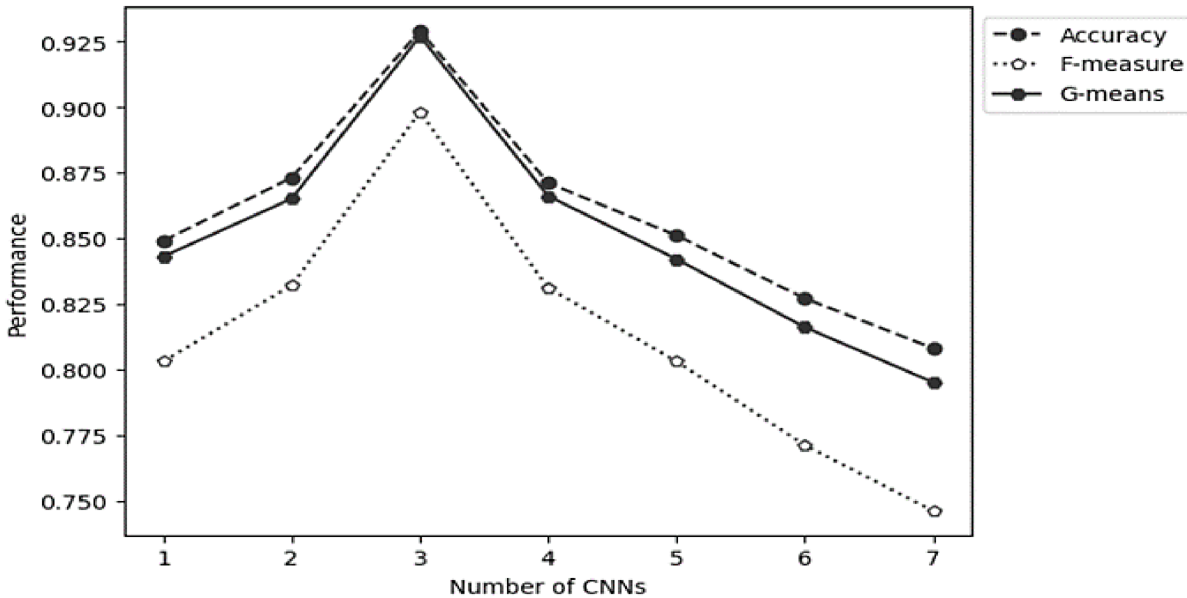

Fig. 6. Plotting the performance indicators of the suggested model while altering the quantity of convolutional feature extraction layers.

D. Discussion

The proposed model signifies a significant advancement in the landscape of anomaly detection methods within the domain of SHM. By incorporating dilated convolutional and DE and RL techniques, this model showcases impressive predictive accuracy. These strides in technological innovation are especially pertinent considering the current challenges that the field of civil infrastructure encounters on a global scale.

However, it is essential to subject the model to critical examination within a broader context of its applicability. While the initial results are promising, they are inherently tied to data originating from a singular architectural marvel – a long-span cable-stayed bridge situated in China. While an in-depth focus on a specific dataset can yield valuable insights, it also presents the potential risk of confining the model to a narrow scope. Civil engineering marvels around the world encompass an immense range – from complex metro rail networks navigating urban mazes to towering skyscrapers reaching for the skies. Each of these structures is the culmination of distinct combinations of design, materials, and environmental factors, leading to unique challenges in structural health. For example, a dam nestled within mountainous terrain would encounter vastly different issues compared to a highway bridge spanning a saline estuary. Each structure reacts to external influences in a nuanced manner, whether it is the ceaseless battering of waves, vehicular loads, or the immense pressure of contained water. Therefore, while the anomalies identified in the Chinese bridge dataset offer invaluable insights, they might only scratch the surface of potential structural concerns when considering the full spectrum of potential issues.

Furthermore, alongside the diversity in structures, the environments in which they exist introduce an additional layer of complexity. The health of a structure isn't solely a reflection of its construction but also a result of its interactions with the environment [37]. From corrosion due to saline exposure to vibrations induced by seismic activities, the array of external stressors is extensive. This raises legitimate concerns about whether the proposed model, primarily trained on the Chinese bridge dataset, can seamlessly adapt to the myriad challenges that structures worldwide encounter. To address these concerns, several solutions can be implemented:

- Diverse Data Collection: Expanding the training dataset to include data from structures in different environmental conditions and geographic locations.

This would enhance the model's ability to generalize across a wide range of scenarios [38].

- Environmental Conditioning: Integrating environmental factors into the model, allowing it to learn how different environmental conditions affect structural health. This could involve adding parameters that account for local climate, pollution levels, and other relevant environmental data.

- Transfer Learning: Applying transfer learning techniques to adapt the model trained on the Chinese bridge dataset to other structures. This approach involves fine-tuning the model with smaller datasets from different structures, enabling it to adjust to new environments with minimal data.

- Regular Model Updates: Continuously updating the model with new data collected from various structures over time. This would ensure that the model stays relevant and effective in predicting structural health under changing environmental conditions [3].

- Hybrid Modeling Approaches: Combining the strengths of different modeling techniques, such as physics-based models and data-driven models [39]. This hybrid approach can leverage the accuracy of physics-based models in well-understood scenarios and the flexibility of data-driven models in complex, variable conditions.

- Real-time Environmental Monitoring: Integrating real-time environmental monitoring systems to provide continuous input to the model. This would allow the model to adjust its predictions based on current environmental conditions.

- Stress Testing and Simulations: Conducting stress tests and simulations under various environmental conditions to validate and improve the model's accuracy in different scenarios.

As the field of civil engineering advances, embracing new materials and groundbreaking construction techniques, the characteristics of potential structural irregularities are likely to transform. A cutting-edge SHM system must be proficient in identifying established problems and adept at signaling new, unexplored issues [40]. This capacity forms a crucial benchmark that the proposed model must meet. The implications are significant; failing to detect a key anomaly can result in catastrophic events, loss of human lives, and severe economic consequences. To enhance the proposed model's capability in this dynamic field, several approaches can be considered:

- Incorporation of Advanced Learning Algorithms: Utilizing machine learning and artificial intelligence algorithms that are capable of identifying patterns and anomalies not only from past data but also adapting to new trends. Techniques like unsupervised learning or deep learning can be particularly effective in recognizing unforeseen issues.

- Continuous Model Updating and Training: Regularly updating the model with the latest data from ongoing construction projects and newly developed materials. This will ensure that the model stays current and can recognize anomalies associated with new construction methodologies.

- Collaborative Data Sharing: Establishing a collaborative network with other civil engineering projects and research institutions for sharing data and insights. This collective approach can significantly broaden the spectrum of scenarios the model is exposed to, enhancing its ability to identify a wide range of anomalies.

- Predictive Analytics: Incorporating predictive analytics to forecast potential structural issues based on current trends and construction practices. This proactive approach can help in early identification and prevention of structural failures.

- Cross-Disciplinary Integration: Integrating knowledge from other fields such as materials science, meteorology, and environmental engineering. This interdisciplinary approach can provide a more comprehensive understanding of how various factors might contribute to new types of structural anomalies.

- Regular Sensitivity Analysis and Testing: Performing sensitivity analyses and stress tests under a variety of conditions to evaluate the model's effectiveness in detecting anomalies in different materials and construction methods.

- Expert Involvement and Feedback Loops: Engaging industry experts in regular reviews of the model's performance, ensuring that practical, real-world insights are incorporated. Establishing feedback loops can also aid in continuous improvement of the model.

As we chart our course ahead, multiple paths invite investigation. Firstly, testing the proposed model against a range of SHM datasets that include different types of structures, such as high-rise buildings, bridges, tunnels, and historical monuments, could reveal its extensive applicability [41]. Employing transfer learning techniques to adapt pre-existing models to these varied scenarios could be key in rapidly broadening the model's utility without necessitating extensive data gathering from each new structure type. In addition to these approaches, several other strategies could be beneficial:

- Cross-Functional Collaboration: Engaging with experts from different fields within civil engineering and data science to gain insights into specific structural characteristics and data processing techniques. This collaboration could enhance the model's accuracy and relevance across various structures.

- Real-Time Data Integration: Incorporating real-time monitoring data into the model to continually update and refine its predictive capabilities. This could include data from sensors monitoring weather conditions, material fatigue, and other relevant parameters.

- Customizable Model Parameters: Developing the model with customizable parameters that can be adjusted according to the specific requirements of different structures. This flexibility would allow for tailored applications, enhancing the model's effectiveness across diverse structural contexts.

- Scalability and Efficiency Improvements: Optimizing the model for scalability and computational efficiency to handle large datasets and enable its deployment in large-scale projects, such as city-wide infrastructure monitoring.

- Community Engagement and Feedback: Involving community feedback, especially from those who live or work in or near monitored structures, to provide ground-level insights into the model's performance and impact.

- Robust Validation and Testing: Conducting rigorous validation and testing under various conditions and scenarios to ensure the model's reliability and accuracy, particularly in critical and emergency situations.

- Policy and Regulatory Alignment: Ensuring that the model aligns with existing policies, standards, and regulatory requirements related to structural health and safety, to facilitate its acceptance and implementation.

Furthermore, the ever-changing characteristics of civil structures require that our SHM systems adapt and improve constantly. Implementing online learning paradigms in the proposed model would enable it to dynamically adjust to evolving structural health patterns. This could be achieved by continuously feeding the model with live data and allowing it to learn and update its parameters in real-time. The integration of diverse data sources, such as vibrations, strains, temperature changes, acoustic emissions, and even visual data from inspections, would significantly enrich the model's predictive accuracy [42]. Several additional steps can be taken to enhance the model's utility and efficiency:

- Edge Computing Implementation: Developing the model for deployment in edge computing environments where data processing occurs closer to the data source. This reduces latency and can be crucial for timely decision-making, especially in emergency scenarios.

- User-Friendly Interface Development: Creating intuitive user interfaces for the model that enable engineers and maintenance personnel to easily interpret and act upon the data and predictions provided by the system.

- Automated Alert and Reporting System: Integrating an automated system that generates alerts and detailed reports when anomalies are detected, thereby facilitating prompt and informed responses from the relevant authorities or maintenance teams.

- Interoperability with Existing Systems: Ensuring that the model is compatible with existing infrastructure management systems and can be seamlessly integrated into current workflows, enhancing its practicality and adoption.

- Regular Benchmarking and Validation: Regularly comparing the model's performance with other state-of-the-art anomaly detection systems in the field to validate its effectiveness and identify areas for improvement.

- Sustainability and Environmental Impact Assessment: Considering the environmental impact and sustainability of the model, especially in terms of its energy consumption and the materials required for sensor deployment and maintenance.

- Training and Education for Stakeholders: Providing comprehensive training and educational resources for engineers, technicians, and stakeholders to understand and effectively use the model in their operations.

Finally, it is important to delve deeper into the specific mechanisms and algorithms used in state-of-the-art techniques for handling imbalanced datasets in RL [43]. This involves examining different approaches, such as oversampling, under sampling, synthetic data generation, cost-sensitive learning, and novel reward shaping strategies [44, 45]. By contrasting these methods with our own, we can identify unique advantages or shortcomings in both theoretical and practical applications. Investigating how these techniques perform in diverse RL environments, ranging from simulated tasks to real-world applications, will provide a more holistic understanding of their adaptability and robustness. It would also be beneficial to explore the integration of our technique with other advanced machine learning strategies like deep learning, transfer learning, and meta-learning, to enhance its performance in handling imbalanced datasets. Such an in-depth analysis will not only fortify our research but also pave the way for future innovations in the field, fostering a more effective approach to tackling the challenges posed by imbalanced datasets in reinforcement learning [46].

VI. Conclusion

This study introduced a groundbreaking model meticulously crafted to confront the intricate challenges associated with anomaly classification within SHM data. The proposed model harnessed a strategic fusion of dilated convolutional, RL, and DE techniques to achieve a high level of accuracy in its results. At its core, the model utilized a group of CNNs to extract essential feature vectors from input images concurrently. These extracted features were seamlessly integrated into downstream processes, bolstering the model's prowess in identifying complex patterns present in SHM data. The efficacy of the proposed model was rigorously validated through experimentation on an imbalanced dataset obtained from a long-span cable-stayed bridge in China-sourced from the IPC-SHM community. Handling imbalanced datasets poses distinct challenges in training classifiers, as the overrepresented class often exerts a disproportionate influence on the learning process, leading to suboptimal performance for the underrepresented class. To effectively address this concern, a novel approach was employed, integrating RL principles to formulate the training procedure as a series of interconnected decisions. Within this framework, the dataset samples assumed the role of states, while the model operated as the agent, receiving appropriate rewards or penalties based on accurate or incorrect classifications, respectively. This adaptive strategy enabled the model to place a heightened focus on the underrepresented class, thereby enhancing classification outcomes. An innovative contribution to the training methodology was introduced by incorporating a mutation operator grounded in clustering principles within the framework of DE. This approach initiated the BP process by identifying a prominent cluster within the existing DE population. Subsequently, a novel update strategy was implemented to generate potential solutions, adding a layer of sophistication to the training process. The experimental results underscored the superior performance of the proposed model in the detection of multi pattern anomalies within SHM data, showcasing remarkable accuracy. Through the adept amalgamation of dilated convolutional, RL, and DE techniques, the model exhibited its potential as an advanced tool for anomaly detection within SHM systems. This capability is of utmost importance in safeguarding the structural integrity and safety of critical infrastructures, including vital components like bridges.

REFERENCES

[1] J. P. Lynch, C. R. Farrar, and J. E. Michaels, "Structural health monitoring: technological advances to practical implementations [scanning the issue]," Proceedings of the IEEE, vol. 104, no. 8, pp. 1508-1512, 2016.

[2] W.-H. Hu, S. Said, R. G. Rohrmann, Á. Cunha, and J. Teng, "Continuous dynamic monitoring of a prestressed concrete bridge based on strain, inclination and crack measurements over a 14-year span," Structural Health Monitoring, vol. 17, no. 5, pp. 1073-1094, 2018.

[3] Z. Tang, Z. Chen, Y. Bao, and H. Li, "Convolutional neural network‐based data anomaly detection method using multiple information for structural health monitoring," Structural Control and Health Monitoring, vol. 26, no. 1, p. e2296, 2019.

[4] H. Han, W.-Y. Wang, and B.-H. Mao, "Borderline-SMOTE: a new over-sampling method in imbalanced data sets learning," in International conference on intelligent computing, 2005: Springer, pp. 878-887.

[5] I. Mani and I. Zhang, "kNN approach to unbalanced data distributions: a case study involving information extraction," in Proceedings of workshop on learning from imbalanced datasets, 2003, vol. 126: ICML, pp. 1-7.

[6] A. Fernández, S. García, M. Galar, R. C. Prati, B. Krawczyk, and F. Herrera, Learning from imbalanced data sets. Springer, 2018.

[7] S. Wang, W. Liu, J. Wu, L. Cao, Q. Meng, and P. J. Kennedy, "Training deep neural networks on imbalanced data sets," in 2016 international joint conference on neural networks (IJCNN), 2016: IEEE, pp. 4368-4374.

[8] C. Huang, Y. Li, C. C. Loy, and X. Tang, "Learning deep representation for imbalanced classification," in Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 5375-5384.

[9] E. Lin, Q. Chen, and X. Qi, "Deep reinforcement learning for imbalanced classification," Applied Intelligence, vol. 50, pp. 2488-2502, 2020.

[10] S. V. Moravvej, S. J. Mousavirad, D. Oliva, and F. Mohammadi, "A Novel Plagiarism Detection Approach Combining BERT-based Word Embedding, Attention-based LSTMs and an Improved Differential Evolution Algorithm," arXiv preprint arXiv:2305.02374, 2023.

[11] T. Eltaeib and A. Mahmood, "Differential evolution: A survey and analysis," Applied Sciences, vol. 8, no. 10, p. 1945, 2018.

[12] S. Samudra, M. Barbosh, and A. Sadhu, "Machine Learning-Assisted Improved Anomaly Detection for Structural Health Monitoring," Sensors, vol. 23, no. 7, p. 3365, 2023.

[13] Q. Pan, Y. Bao, and H. Li, "Transfer learning-based data anomaly detection for structural health monitoring," Structural Health Monitoring, p. 14759217221142174, 2023.

[14] S. Li, L. Jin, Y. Qiu, M. Zhang, and J. Wang, "Signal anomaly detection of bridge SHM system based on two-stage deep convolutional neural networks," Structural Engineering International, vol. 33, no. 1, pp. 74-83, 2023.

[15] X. Ye, P. Wu, A. Liu, X. Zhan, Z. Wang, and Y. Zhao, "A Deep Learning-based Method for Automatic Abnormal Data Detection: Case Study for Bridge Structural Health Monitoring," International Journal of Structural Stability and Dynamics, p. 2350131, 2023.

[16] D. K. Green and A. Jaspan, "Applied Bayesian Structural Health Monitoring: inclinometer data anomaly detection and forecasting," arXiv preprint arXiv:2307.00305, 2023.

[17] R. Boccagna, M. Bottini, M. Petracca, A. Amelio, and G. Camata, "Unsupervised Deep Learning for Structural Health Monitoring," Big Data and Cognitive Computing, vol. 7, no. 2, p. 99, 2023.

[18] X. Lei, Y. Xia, A. Wang, X. Jian, H. Zhong, and L. Sun, "Mutual information based anomaly detection of monitoring data with attention mechanism and residual learning," Mechanical Systems and Signal Processing, vol. 182, p. 109607, 2023.

[19] Y. Yang and S. Nagarajaiah, "Data compression of structural seismic responses via principled independent component analysis," Journal of Structural Engineering, vol. 140, no. 7, p. 04014032, 2014.

[20] Y. Yang and S. Nagarajaiah, "Robust data transmission and recovery of images by compressed sensing for structural health diagnosis," Structural Control and Health Monitoring, vol. 24, no. 1, p. e1856, 2017.

[21] Y. Yang and S. Nagarajaiah, "Blind denoising of structural vibration responses with outliers via principal component pursuit," Structural Control and Health Monitoring, vol. 21, no. 6, pp. 962-978, 2014.

[22] S. Park, S. Kim, and J.-H. Choi, "Gear fault diagnosis using transmission error and ensemble empirical mode decomposition," Mechanical Systems and Signal Processing, vol. 108, pp. 262-275, 2018.

[23] Y. Bao, J. Li, T. Nagayama, Y. Xu, B. F. Spencer Jr, and H. Li, "The 1st international project competition for structural health monitoring (IPC-SHM, 2020): A summary and benchmark problem," Structural Health Monitoring, vol. 20, no. 4, pp. 2229-2239, 2021.

[24] J. Parra, L. Trujillo, and P. Melin, "Hybrid back-propagation training with evolutionary strategies," Soft Computing, vol. 18, no. 8, pp. 1603-1614, 2014.

[25] K. Deb, "A population-based algorithm-generator for real-parameter optimization," Soft Computing, vol. 9, pp. 236-253, 2005.

[26] V. François-Lavet, P. Henderson, R. Islam, M. G. Bellemare, and J. Pineau, "An introduction to deep reinforcement learning," Foundations and Trends® in Machine Learning, vol. 11, no. 3-4, pp. 219-354, 2018.

[27] M. Bahadori, M. Soltani, M. Soleimani, and M. Bahadori, "Statistical Modeling in Healthcare: Shaping the Future of Medical Research and Healthcare Delivery," in AI and IoT-Based Technologies for Precision Medicine: IGI Global, 2023, pp. 431-446.

[28] S. Danaei et al., "Myocarditis Diagnosis: A Method using Mutual Learning-Based ABC and Reinforcement Learning," in 2022 IEEE 22nd International Symposium on Computational Intelligence and Informatics and 8th IEEE International Conference on Recent Achievements in Mechatronics, Automation, Computer Science and Robotics (CINTI-MACRo), 2022: IEEE, pp. 000265-000270.

[29] E. Barron and H. Ishii, "The Bellman equation for minimizing the maximum cost," NONLINEAR ANAL. THEORY METHODS APPLIC., vol. 13, no. 9, pp. 1067-1090, 1989.

[30] J. Mao, H. Wang, and B. F. Spencer Jr, "Toward data anomaly detection for automated structural health monitoring: Exploiting generative adversarial nets and autoencoders," Structural Health Monitoring, vol. 20, no. 4, pp. 1609-1626, 2021.

[31] M. Zhao, A. Sadhu, and M. Capretz, "Multiclass anomaly detection in imbalanced structural health monitoring data using convolutional neural network," Journal of Infrastructure Preservation and Resilience, vol. 3, no. 1, p. 10, 2022.

[32] Ö. Özdemir and E. B. Sönmez, "Weighted cross-entropy for unbalanced data with application on covid x-ray images," in 2020 Innovations in

Intelligent Systems and Applications Conference (ASYU), 2020: IEEE, pp. 1-6.

[33] F. Huang, J. Li, and X. Zhu, "Balanced Symmetric Cross Entropy for Large Scale Imbalanced and Noisy Data," arXiv preprint arXiv:2007.01618, 2020.

[34] X. Li, X. Sun, Y. Meng, J. Liang, F. Wu, and J. Li, "Dice loss for data-imbalanced NLP tasks," arXiv preprint arXiv:1911.02855, 2019.

[35] S. S. M. Salehi, D. Erdogmus, and A. Gholipour, "Tversky loss function for image segmentation using 3D fully convolutional deep networks," in Machine Learning in Medical Imaging: 8th International Workshop, MLMI 2017, Held in Conjunction with MICCAI 2017, Quebec City, QC, Canada, September 10, 2017, Proceedings 8, 2017: Springer, pp. 379-387.

[36] S. A. Taghanaki et al., "Combo loss: Handling input and output imbalance in multi-organ segmentation," Computerized Medical Imaging and Graphics, vol. 75, pp. 24-33, 2019.

[37] A. Moallemi, A. Burrello, D. Brunelli, and L. Benini, "Model-based vs. data-driven approaches for anomaly detection in structural health monitoring: A case study," in 2021 IEEE International Instrumentation and Measurement Technology Conference (I2MTC), 2021: IEEE, pp. 1-6.

[38] Y.-M. Zhang, H. Wang, H.-P. Wan, J.-X. Mao, and Y.-C. Xu, "Anomaly detection of structural health monitoring data using the maximum likelihood estimation-based Bayesian dynamic linear model," Structural Health Monitoring, vol. 20, no. 6, pp. 2936-2952, 2021.

[39] C. Bigoni, "Numerical methods for structural anomaly detection using model order reduction and data-driven techniques," EPFL, 2020.

[40] Y. Bao, Z. Tang, H. Li, and Y. Zhang, "Computer vision and deep learning–based data anomaly detection method for structural health monitoring," Structural Health Monitoring, vol. 18, no. 2, pp. 401-421, 2019.

[41] Y. Zhang, Z. Tang, and R. Yang, "Data anomaly detection for structural health monitoring by multi-view representation based on local binary patterns," Measurement, vol. 202, p. 111804, 2022.

[42] X. Xu et al., "Anomaly detection for large span bridges during operational phase using structural health monitoring data," Smart Materials and Structures, vol. 29, no. 4, p. 045029, 2020.

[43] S. Susan and A. Kumar, "The balancing trick: Optimized sampling of imbalanced datasets—A brief survey of the recent State of the Art," Engineering Reports, vol. 3, no. 4, p. e12298, 2021.

[44] K. M. Hasib et al., "A survey of methods for managing the classification and solution of data imbalance problem," arXiv preprint arXiv:2012.11870, 2020.

[45] D. Ramyachitra and P. Manikandan, "Imbalanced dataset classification and solutions: a review," International Journal of Computing and Business Research (IJCBR), vol. 5, no. 4, pp. 1-29, 2014.

[46] J. M. Johnson and T. M. Khoshgoftaar, "Survey on deep learning with class imbalance," Journal of Big Data, vol. 6, no. 1, pp. 1-54, 2019.