BER the international journal of bridge engineering management and research

An Adversarial-Based Method for Knowledge Transfer Across Bridges

Valentina Giglioni*; Ilaria Venanzi; and Filippo Ubertini

Submitted: 07 June 2025 **Accepted:** 23 June 2025 **Publication date:** 10 July 2025

DOI: 10.70465/ber.v2i3.42

Abstract: Data-driven machine learning methods for bridge health assessment often face limitations due to the scarcity of labeled data, particularly related to damaged conditions, which are often difficult, expensive, or even impossible to obtain in practice. To address this challenge, transfer learning offers a promising solution by leveraging knowledge gained from one structure (or domain), with sufficient labeled data, to identify damage on a different but related structure where labeled data are limited or unavailable. In this study, we propose a domain adversarial neural network-based method to enhance damage classification performance across different bridge structures. Using the natural frequencies collected over long-term monitoring campaigns, the machine learning model is trained with an adversarial strategy to learn damage-sensitive and domain-invariant features for improving generalization to new structures with minimal additional data collection. To validate the methodology, transfer learning results are analyzed by first considering two post-tensioned concrete bridges, the Z24 bridge and the S101 bridge, and afterwards their finite element models, where different damage scenarios in terms of localization and severity are simulated. Furthermore, the performance of the proposed method is compared with different transfer learning approaches previously applied in the same structural health monitoring context. The improvement of damage classification results via transfer learning highlights the potential of domain-adversarial learning to advance scalable monitoring strategies for bridge networks.

Author keywords: Domain-adversarial neural network; SHM; transfer learning; bridge network

Introduction

Over recent decades, vibration-based structural health monitoring (SHM) has gained widespread attention as a tool to manage and preserve aging infrastructure, which is increasingly exposed to material degradation and evolving operational demands.[1,2] Advancements in sensing technologies and computational capabilities have accelerated the adoption of data-driven approaches, particularly in the use of machine learning (ML) algorithms, to analyze dynamic monitoring data and infer informative trends, anomalies, or specific patterns that provide insights for predicting structural assessment.[3,4] By applying supervised or unsupervised learning, ML models can automate multiple SHM tasks, including damage detection, localization, and quantification during long-term monitoring campaigns. Applications of different ML methods are described in Wedel et al.[5] to detect and compensate for sensor faults and forecast the behavior of real monitored bridges.

Despite their potential, ML models typically require large amounts of labeled data to generalize well. In practice, such data, especially from damaged conditions, are costly, scarce, and often infeasible to obtain. Moreover, most conventional ML algorithms assume that training and testing data come from the same distribution. This assumption breaks down when models trained on one structure are applied to another with different properties, limiting the scalability of SHM solutions.

To overcome these challenges, transfer learning (TL) has emerged as a promising approach,[6] enabling the transfer of diagnostic knowledge from a fully labeled source structure to a target structure with limited or no labeled data. By leveraging labels from one or multiple similar source structures, TL can improve performance and reduce the need for a structure-specific damage-identification algorithm, addressing a key bottleneck in real-world SHM deployment. Within the bridge monitoring framework, transductive methods such as domain adaptation (DA) are often more feasible, since target data is typically accessible, even if unlabeled. The core idea is to find a mapping that mitigates the distribution discrepancy between two distinct domains. Figueiredo et al.[7] applied transfer component analysis to match health-state data referring to different environmental conditions between the Z24 bridge and the finite element model (FEM) of the same bridge, without transferring damage labels. A similar approach was carried out in Yano et al.,[8] where

*Corresponding Author: Valentina Giglioni.
Email: valentina.giglioni@unipg.it
Department of Civil and Environmental Engineering, University of Perugia, Via G. Duranti 93, Perugia 06125, Italy

Discussion period open till six months from the publication date. Please submit separate discussion for each individual paper. This paper is a part of the Vol. 2 of the International Journal of Bridge Engineering, Management and Research (© BER), ISSN 3065-0569.

the application of TL was devoted to novelty detection, aimed at identifying deviations from the normal state. Poole et al.[9] proposed a statistic alignment method, known as normal condition alignment (NCA), to align the lower-order statistics of source and target domains in the original feature space, proving it to be robust for poor and limited datasets. The effectiveness of a combined use of NCA and joint domain adaptation (JDA)[10] was evaluated in Giglioni et al.,[11] aiming to enable supervised damage detection and classification across the real Z24 and S101 bridges and their corresponding FEMs, as well as in Giglioni et al.,[12] which presents a physical benchmark via lab-scale bridges and the results of single-source and multisource DA. In the realm of DA, the domain adversarial neural network (DANN) theory is recently being developed to address TL-guided SHM. DANNs leverage adversarial learning to train a feature extractor that simultaneously maximizes label classification performance and minimizes domain discrepancy. This is accomplished through the integration of a gradient reversal layer (GRL), which encourages the extracted features to be predictive of class labels while being indistinguishable across domains. Beyond several works dedicated to bearing fault diagnosis[13–15] and damage identification in building structures,[16] applications of these techniques in the field of bridge SHM are still limited. Among these, Li et al.[17] implemented a bridge damage detection method using moving load-induced displacement response, while Li et al.[18] presented a conditional adversarial DA method tested on beam models.

By promoting domain invariance without compromising task-relevant information, DANN is here adopted to perform unsupervised DA via the transfer of damage-state knowledge from the source to the completely unlabeled target domain. As main contributions, the procedure is validated on both full-scale real infrastructures, the Z24 and the S101 bridges, and considering numerical models with multiple simulated damage scenarios, enabling performance assessment across varying locations and damage severity, as well as multi-class damage identification. Furthermore, this study represents an advancement of the work presented in Giglioni et al.,[11] since the DANN-focused methodology is applied to the same case studies and compares the results associated with different DA approaches, pointing out the advantages and strengths of adversarial networks, particularly regarding computational efficiency and effectiveness.

As bridge systems grow increasingly complex and the demand for sustainable infrastructure intensifies, TL is poised to play a pivotal role in advancing future monitoring and maintenance strategies. This paper presents promising results from the proposed DANN-based approach, highlighting its potential for effective bridge monitoring even with no prior information about bridges subjected to new monitoring campaigns. At the same time, key challenges and open questions that warrant further investigation are discussed and pointed out. The structure of the paper is as follows: Section 2 outlines the fundamental principles of DANNs and the adopted methodology; Section 3 introduces the two benchmark bridges and the corresponding FEMs; Section 4 presents the results of damage detection and classification following DA, showing a comparison with previous results; and Section 5 concludes the paper and summarizes the key findings.

TL Methodology via DANNs

DA

A brief introduction to TL and its sub-discipline, DA, is provided hereafter. A domain $\mathcal{D} = \{\mathcal{X}, p(X)\}$ can be represented by a feature space $\mathcal{X}$ and a marginal probability distribution $p(X)$, where $X = \{x_i\}_{i=n}^{N}$ is a finite sample set from $\mathcal{X}$. A task $\mathcal{T} = \{\mathcal{Y}, f(\cdot)\}$ is identified by a label space $\mathcal{Y}$ and a predictive function $f(\cdot)$. TL aims to improve the predictive function of a target domain $\mathcal{D}_t$ using knowledge extracted from $\mathcal{D}_s$ and $\mathcal{T}_s$, assuming $\mathcal{D}_s \neq \mathcal{D}_t$ and/or $\mathcal{T}_s \neq \mathcal{T}_t$.

To reduce the gap between feature spaces of source and target domains in terms of proper statistical distances, DA is introduced to minimize the difference in marginal distributions $p(X_s) \neq p(X_t)$, or conditional distributions $p(Y_s|X_s) \neq p(Y_t|X_t)$, or both, but assuming that the feature and label spaces are equal, i.e., $\mathcal{X}_s = \mathcal{X}_t$ and $\mathcal{Y}_s = \mathcal{Y}_t$. The way to perform a specific task in the source domain can therefore be leveraged to predict unknown data pertaining to the target domain.

Proposed DA framework

The architecture of the proposed network is illustrated in Fig. 1.

The process begins with the collection of data representing both undamaged and damaged conditions from a fully labeled source domain. Operational modal analysis is carried out to continuously extract the set of natural frequencies characterizing the two-dimensional input matrix, where each row represents an acquired sample and each column contains the corresponding natural frequencies. With the aim of representing the most common situation in real-world monitoring scenarios, the target domain (i.e., the bridge to be assessed) is assumed to be entirely unlabeled.

The input to the model consists of natural frequency time histories, selected as damage-sensitive features, continuously recorded via long-term, continuous monitoring. The network includes three main components, a feature extractor, a label predictor, and a domain classifier. These blocks are trained jointly to learn latent representations that should be (i) discriminative for accurately identifying various health states and (ii) domain-invariant to enable effective generalization across different domains. Specifically, the feature extractor $G_f(\cdot, \theta_f)$ acts as a shared encoder that processes input data from both domains to capture the essential information needed for the other two components. In line with the adversarial learning paradigm, the label predictor $G_y(\cdot, \theta_y)$ utilizes the extracted features to learn class predictions based on the labeled input data from the source domain. Simultaneously, the domain classifier $G_d(\cdot, \theta_d)$ attempts to determine the origin (source or target) of the input data. The learnable parameters, i.e., the weights of the feature extractor, the label

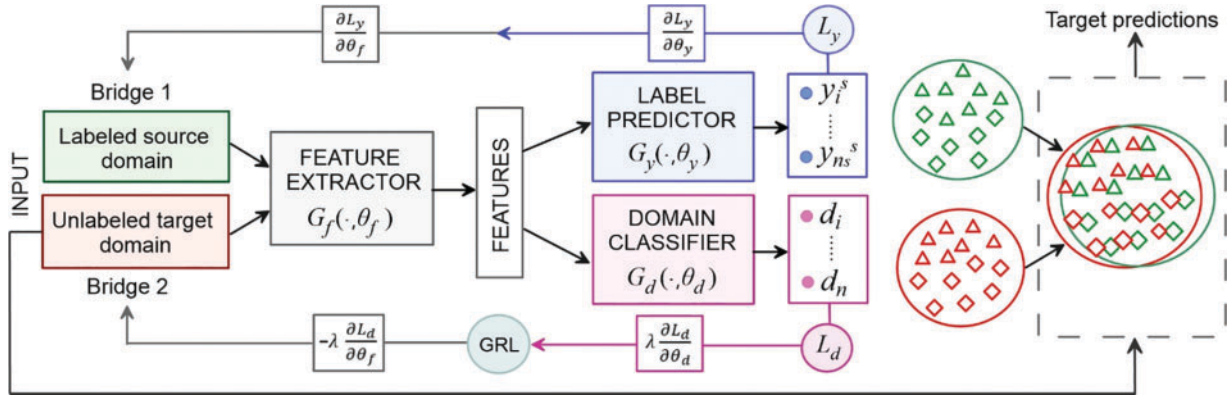

Figure 1. Flowchart of the DANN-based methodology

predictor, and the domain classifier, are indicated as θ_f, θ_y, and θ_d. To extract domain-invariant features and confuse the domain classifier, while a task-specific loss $L_y(\theta_f, \theta_y)$ is minimized, a domain-specific loss $L_d(\theta_f, \theta_d)$ is maximized. A key element enabling this adversarial setup is the GRL, placed between the feature extractor and the domain classifier. During the forward pass, the GRL acts as a transparent layer, passing features unaltered. However, during backpropagation, it multiplies the gradients flowing into the feature extractor by a negative constant ($-\lambda$), thereby reversing the optimization direction (since the domain classifier itself tries to minimize $L_d(\theta_f, \theta_d)$). The label predictor and domain classification losses are shown in Eqs. (1) and (2), respectively, where L_{CE} is the cross-entropy loss, x_i^s is the ith source input (i.e., the natural frequency) with the associated label y_i^s, $d_i \in \{0, 1\}$ is the domain label (0 = source, 1 = target), $x_i \in \{x^s, x^t\}$, while n_s and n are, respectively, the total number of labeled source samples and the total number of domain samples in the current batch.

$$L_y\left(\theta_f, \theta_y\right) = \frac{1}{\mathrm{n_s}} \sum\nolimits_{i=1}^{n_s} L_{CE}\left(G_y\left(G_f\left(x_i^s\right)\right), y_i^s\right) \tag{1}$$

$$L_d\left(\theta_f, \theta_d\right) = \frac{1}{\mathrm{n}} \sum\nolimits_{j=1}^{n} L_{CE}\left(G_d\left(G_f\left(x_i\right)\right), d_i\right) \tag{2}$$

The overall objective function is provided in Eq. (3):

$$\min_{\theta_f,\theta_y} \max_{\theta_d} (L_y\left(\theta_f, \theta_y\right) - \lambda L_d\left(\theta_f, \theta_d\right) \tag{3}$$

Working in the unsupervised learning environment, the implemented DANN is employed to predict the unknown target labels y_i^t by finding a mapping between data distributions of a source and a target bridge, given prior knowledge of source labels y_i^s and all domain labels d_i.

Application Case Studies

In this paper, the TL procedure is validated by transferring and identifying (i) two different damages occurring in real bridges and (ii) several damages, divided into groups, simulated in the calibrated FEMs, with the aim of investigating variation in terms of damage localization, severity, and extent.

Z24 bridge and S101 bridge: structural descriptions and monitoring campaigns

The Z24 bridge was a post-tensioned reinforced concrete (RC) structure built in Switzerland in 1963 and dismantled in late 1998 to allow for the construction of a wider side span. The deck is characterized by a twin-box girder cross-section, with a main span of 30 m and two side spans of 14 m. The overall width was 8.6 m, and the structural height was 1.01 m (see Fig. 2).

Two concrete piers with rectangular cross-sections (0.4 m × 0.32 m × 6 m) were placed at each end span and rigidly connected to the girder. The bridge was continuously monitored from November 1997 to September 1998 using eight accelerometers, which recorded measurements over 10-min intervals each hour. Additional environmental sensors captured measurements before and after each dynamic recording. At the end of summer in 1998, a series of progressive damage tests were artificially applied to the structure to provide an extensive dataset and study the evolution of dynamic properties throughout the damage-induction phase. For the purpose of this work, two macro-categories of damage are considered, including the lowering of one pier and the cutting of tendons along the deck, respectively, named D1 and D2. Further details on the monitoring setup and the related damage detection results can be found in Maeck et al.[19].

Similarly, the S101 bridge was a post-tensioned RC structure in Austria, dating back to the early 1960s,[20] composed of a 32-m main span and two 12-m side spans (Fig. 2). The deck cross-section was defined by a double-webbed T-beam, 7.2 m wide, with the height ranging from 0.9 m at mid-span to 1.7 m above the piers. Prior to its demolition, a short monitoring campaign was conducted in December 2008. The bridge was instrumented with 45 accelerometers that collected data for 5 min each hour at a 500 Hz sampling frequency. Three out of the four monitoring days were specifically dedicated to damage-data acquisition. Since ambient temperatures consistently ranged between −2°C and 2°C, environmental influences on modal characteristics were considered negligible. Detailed descriptions of the SHM setup and test protocols for S101 can be found in Döhler et al.[21].

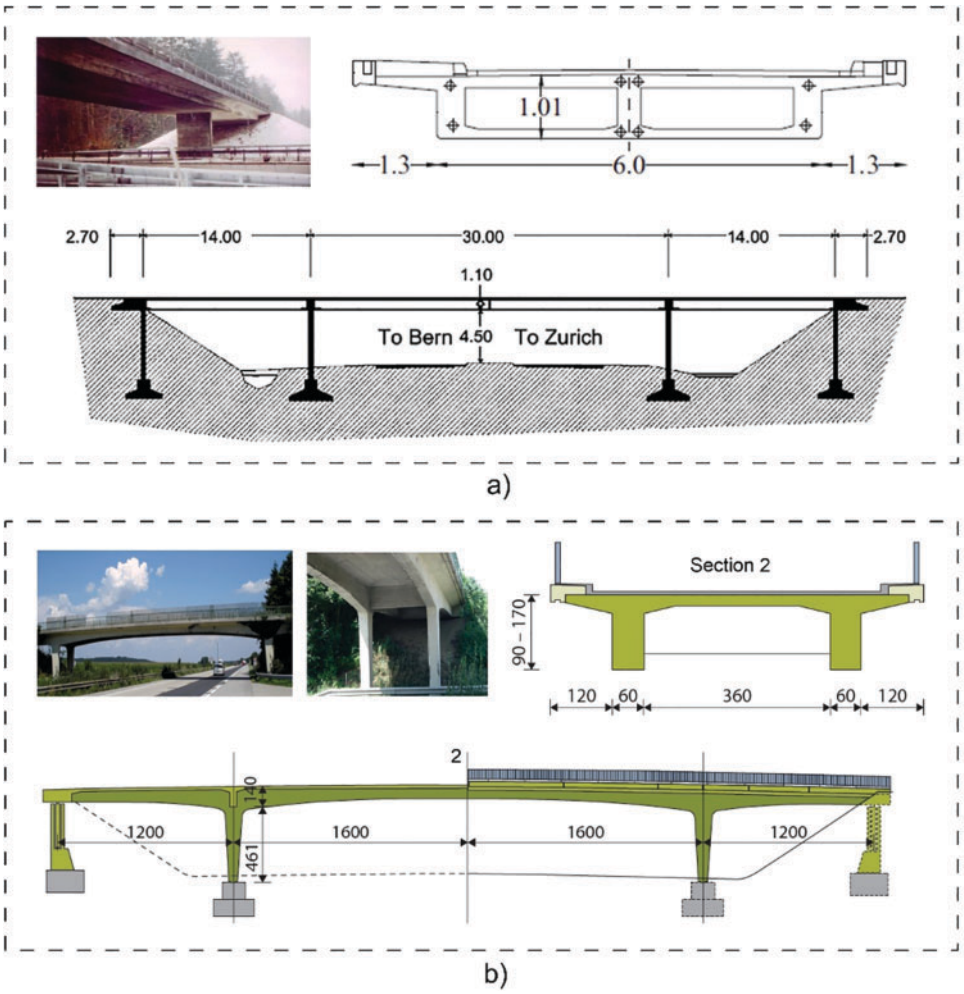

Figure 2. General views and deck cross-sections of the Z24 bridge (a) and the S101 bridge (b)

Table 1. Progressive damage scenarios carried out in the Z24 bridge and the S101 bridge

Class	Z24 bridge	S101 bridge
H	Undamaged state	Undamaged state
D1	Lowering of the pier, 20 mm	North-western pier cut through
	Lowering of the pier, 40 mm	Lowering of the pier, 1 cm
	Lowering of the pier, 80 mm	Lowering of the pier, 2 cm
	Lowering of the pier, 95 mm	Lowering of the pier, 3 cm
	Tilt of foundation	—
D2	Failure of concrete hinge	1st tendon intersected
	Failure of 2 anchor heads	2nd tendon intersected
	Failure of 4 anchor heads	3rd tendon intersected
	Rupture of 2 out of 16 tendons	4th tendon partially intersected
	Rupture of 4 out of 16 tendons	—
	Rupture of 6 out of 16 tendons	—

Note that, although the duration of the damage phases differed, both bridges were subjected to similar damage scenarios, named as D1 and D2, whose descriptions are provided in Table 1.

Therefore, this first case study is devoted to validating the DANN approach by using knowledge from the Z24 bridge to classify the unknown, limited data of the S101 bridge. The dynamic characterization of the two bridges, details of which are omitted as they fall outside the scope of this work, is achieved via operational modal analysis

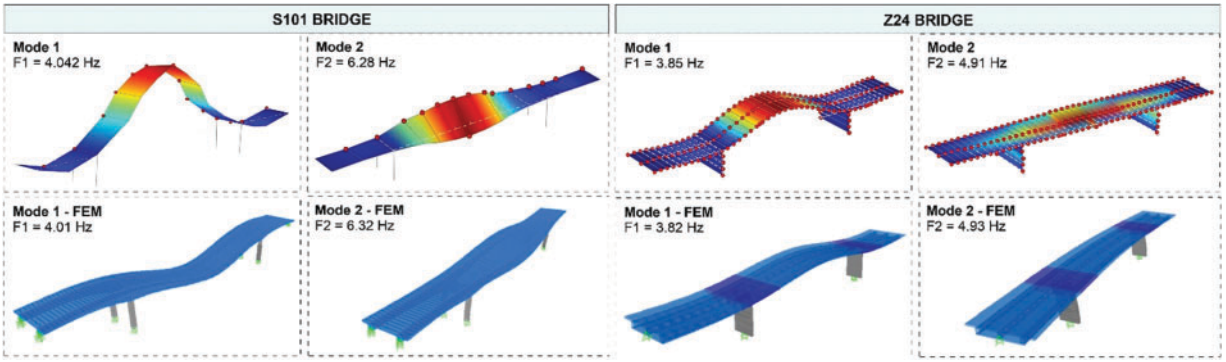

Figure 3. Mode shapes of the selected natural frequencies for the Z24 bridge and the S101 bridge

Table 2. Data sets size

Class	Z24 bridge	S101 bridge
H	(4115, 2)	(32, 2)
D1	(175, 2)	(14, 2)
D2	(13, 2)	(29, 2)

(OMA) and automated frequency tracking using the MOSS software, as implemented and illustrated in García-Macías et al.[22]. As shown in Fig. 3, despite different absolute values in natural frequencies, clear similarities in modal responses can be highlighted. Hence, the first two natural frequencies, respectively, describing a bending and a lateral/torsional mode common to both structures, are here adopted as input features. Considering this, data set information regarding the size of each class is summarized in Table 2.

FEM of the Z24 bridge and S101 bridge: structural descriptions and damage scenarios

The FEM of the Z24 bridge is built and calibrated in the SAP environment, using beam finite elements for the deck and the piers. As indicated in Fig. 3, the section above the supporting piers is modeled with a thicker slab because of the presence of higher stiffness values at the end of the main span.[23] In addition to the dead load, a distributed vertical mass and 20 concentrated masses applied in both the transverse and vertical directions are introduced to simulate external loads such as barriers, guardrails, and sidewalks. The concrete elastic modulus (Ec), the mass values, and the thickness of the girder plate are treated as calibration parameters to minimize the discrepancy between the experimentally identified natural frequencies and those estimated through FEM-based modal analysis.

While the four rectangular piers of the S101 bridge are modeled using beam elements, the deck is discretized into shell elements to represent two longitudinal rectangular beams with variable height, the deck slab, and two transverse rectangular beams located above the piers. The calibration process involves adjusting the mass values and the elastic modulus of concrete to minimize discrepancies in natural frequencies. Detailed calibration results are presented in Giglioni et al.[11]. Fig. 3 illustrates the mode shapes associated with the selected input frequencies: the symmetric bending mode (Mode 1) and the lateral/torsional mode (Mode 2).

Both models are utilized to perform modal analysis and to extract natural frequencies under normal conditions, assuming a variation of the elastic modulus within ±1.5%. To simulate common damage scenarios typically affecting bridge infrastructure, several damage classes are introduced into the models, as illustrated in Fig. 4.

Damage scenarios d_0, d_1, and d_2, correspond to reductions of approximately 10%, 15%, and 25%, respectively, in the elastic modulus over a 2-m section at mid-span, an area subjected to the highest positive bending moments. A more extensive damage scenario is labeled d_3, involving a ∼10% reduction in elastic modulus over a 6-m segment along the middle span. The final scenario d_4 simulates a ∼40% reduction in elastic modulus near the deck-to-pier connection, where a plastic hinge is expected to form over a length equal to 1.5 times the deck cross-section height. To ensure the damage scenarios are realistic, the selected reductions in elastic modulus follow Eurocode guidelines, which state that cracking can result in stiffness reductions of up to 50%.[24] The number of samples extracted from both domains is assumed to be equal, with 100 samples representing the undamaged condition (H) and 60 samples for each simulated damage class.

TL Results Using the DANN Approach

The effectiveness of the DANN procedure is demonstrated in addressing three different tasks, illustrated as follows:

- Task 1 (T1): Knowledge transfer from the Z24 bridge to the S101 bridge, to classify the artificially induced damages.
- Task 2 (T2): Knowledge transfer between the FEMs of the two real bridges by considering a combination of damage scenarios including d_1 and d_4.
- Task 3 (T3): Knowledge transfer between the FEMs of the two real bridges by considering a combination of damage scenarios including d_0, d_2, and d_3.

While the transfer in T2 and T3 is bidirectional, resulting in the sub-tasks T2a, T2b, T3a, and T3b (see Table 3), the

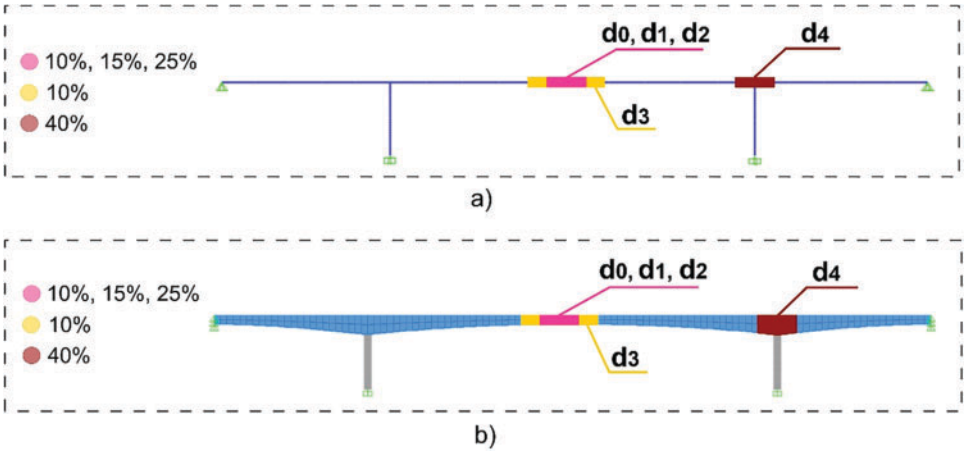

Figure 4. Damage simulations in the Z24 bridge (a) and the S101 bridge (b) FEMs: d_0, d_1, d_2, and d_3 represent, respectively, the damage applied to the middle span, varying in severity and extent; d_4 refers to the elastic modulus reduction at the deck-to-pier connection

model adopted for T1 uses the information extracted from the highly populated Z24 bridge dataset (source domain) to infer diagnosis in the S101 bridge (target domain). Working with FEMs, TL capabilities are evaluated by assuming different damage locations (T2), and the same location with different damage severity and extent (T3).

The natural frequencies selected to describe both real bridges are plotted in Fig. 5a, where source and target distributions are clearly shifted in the original feature space, hampering the generalization of the classifier, but exhibiting similar spatial correlations toward damage classes.

In fact, the feature extractor and the label classifier were associated with a poor damage identification performance before domain alignment, i.e., an F1 score equal to 0.43. To perform DA, the DANN-based methodology is validated and afterwards compared with previously applied techniques, such as NCA and the combination of NCA with JDA. In this case, NCA is adopted for normalization before training the DANN. Tailored networks are implemented for each task, since the calibration process strictly depends on the input data distributions. As a common characteristic, each DANN component is characterized by multilayer perceptron (MLP) networks, connecting the dense layers to the next ones. The input layer is passed through the feature extractor to generate latent representations, which are then used by the label predictor to perform classification. The GRL, sitting between the feature extractor and the domain classifier, reverses the gradient direction during backpropagation via the multiplication by a negative scalar λ, selected equal to 1 in this work following the original DANN formulation and facilitating the training process. Sparse categorical cross-entropy loss and *relu* activation functions are adopted for both G_y and G_d. To take into account the dataset imbalance, higher weights are assigned to underrepresented classes during training by following Eq. (4):

$$\text{weight}_i = \frac{n_{\text{samples}}}{n_{\text{classes}} \times n_i} \tag{4}$$

where n_i is the number of samples in the ith class.

The model parameters that vary across tasks are manually tuned to maximize model accuracy and are listed in Table 3, including the number and size of the layers of each network block, the number of epochs, the batch size, and the learning rate.

Compared with the original feature distribution shown in Fig. 5a, the new domain-invariant space exhibits health-state clusters that are shared across source and target domains, indicating a substantial reduction in the initial domain shift. The transformed features are visualized using t-distributed stochastic neighbor embedding (t-SNE),[25] which models pairwise similarities by converting Euclidean distances into conditional probabilities. As a result, high-dimensional features are effectively mapped into a two-dimensional space, represented by t-SNE dimension 1 versus t-SNE dimension 2 (Fig. 5b). The application of the implemented DANN leads to an F1 score equal to 0.84, considering the presence of some misalignment within the D1 class. It is specified that

Table 3. Model architecture and parameters varying tasks

Parameters	T1	T2	T3
Size $L1_{G_f}$	356	256	356
Size $L2_{G_f}$	116	56	126
Size $L3_{G_f}$	66	25	86
Size $L4_{G_f}$	15	—	25
Size $L1_{G_y}$	3	3	4
Size $L1_{G_d}$	370	356	356
Size $L2_{G_d}$	150	126	126
Size $L3_{G_d}$	57	16	56
Size $L4_{G_d}$	16	–	26
Size $L5_{G_d}$	2	2	2
Epoch	120	100	120
Batch size	64	8	8
Learn. rate	0.001	0.0001	0.0001

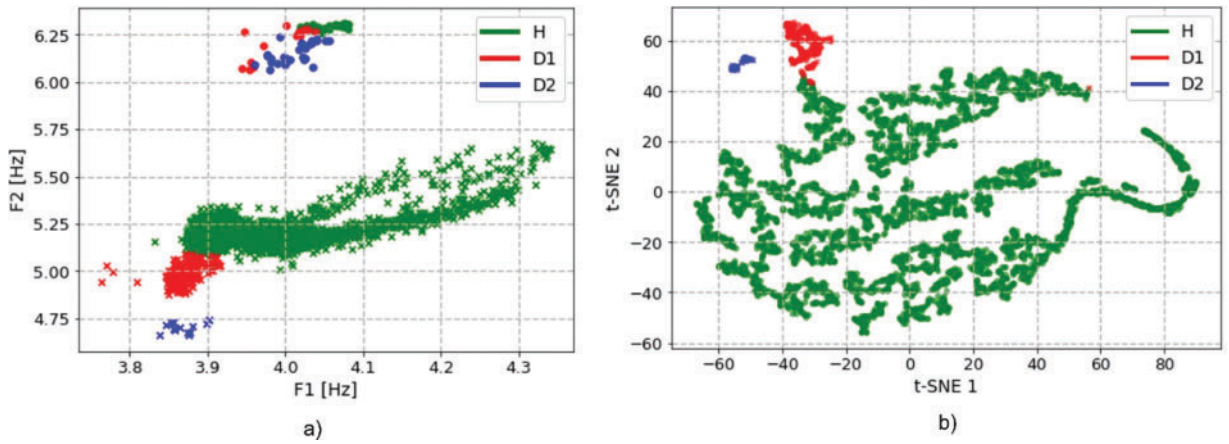

Figure 5. Features distribution of the Z24 bridge (labelled with "×") and S101 bridge (labelled with "◯") before (a) and after (b) DA

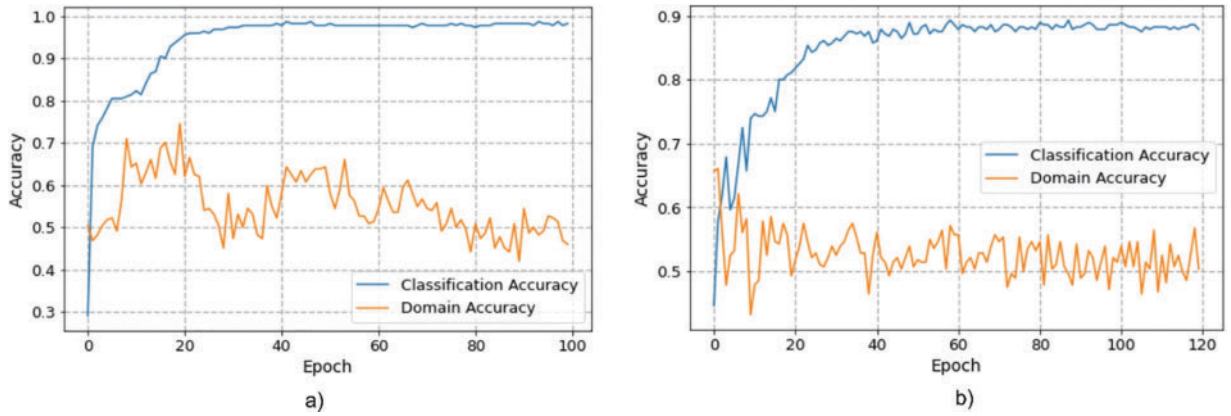

Figure 6. Training process results when transferring knowledge from the S101 FEM bridge to the Z24 FEM bridge to address T2b (a) and T3b (b)

the network model is trained using the entire source dataset and then evaluated on the entire target dataset.

The same approach, but with different model architecture and for different purposes, is tested for the second and third tasks, alternatively assuming the Z24 FEM and the S101 FEM as source models. The original feature spaces are plotted in Figs. 7a and 7b, which include, respectively, the damage scenarios pertaining to T2 and T3. It is easy to underline the gap between source and target distributions and the need to align the data for making TL effective. Following the DANN framework, the time-histories of the frequencies extracted from both FEMs are fed into the networks as model inputs. The goal of the adversarial learning via GRL is to maximize the label predictor accuracy while confusing the domain classifier. This concept is visualized in Fig. 6, showing the opposite behavior of G_y and G_d throughout the epochs, where classification accuracy steadily improves, approaching 1, while domain accuracy decreases until oscillating around 0.5. Precisely, the training processes associated with T2b and T3b are, respectively, reported in Figs. 6a and 6b, when knowledge is transferred from the S101 FEM to the Z24 FEM bridge. To avoid redundancy, the results obtained by inverting source and target domains are omitted due to the exhibited comparable trends.

After unsupervised learning, the DANN is trained and adopted to predict the classes of the completely unlabeled target domain. The resulting latent space, visible in Fig. 7c, demonstrates the capability to make inferences on the S101 FEM by distinguishing between the undamaged state and the damage along the mid-span (d_1) and the pier-to-deck node (d_4), leading to an F1 score equal to 0.97 for the T2a task. The effective alignment is confirmed by the confusion matrix shown in Fig. 8b. Similar results, with an F1 score of 0.98, are also obtained when using the Z24 FEM as the source domain (Fig. 8a). Therefore, the generation of effective latent features via DANN, revealing shared clusters across domains, enables better generalization of the label predictor. The domain-invariant nature of such features is explained by an increase in label predictor accuracy, accompanied by a parallel decrease in the domain classifier performance. Specifically, the domain accuracy drops to values in the range of 0.3–0.5 when evaluating the model on unknown target data, considering all the presented tasks.

Moving to T3, the improvement provided by the TL approach is noticeable when comparing the original feature representation of Fig. 7b with the transformed quantities

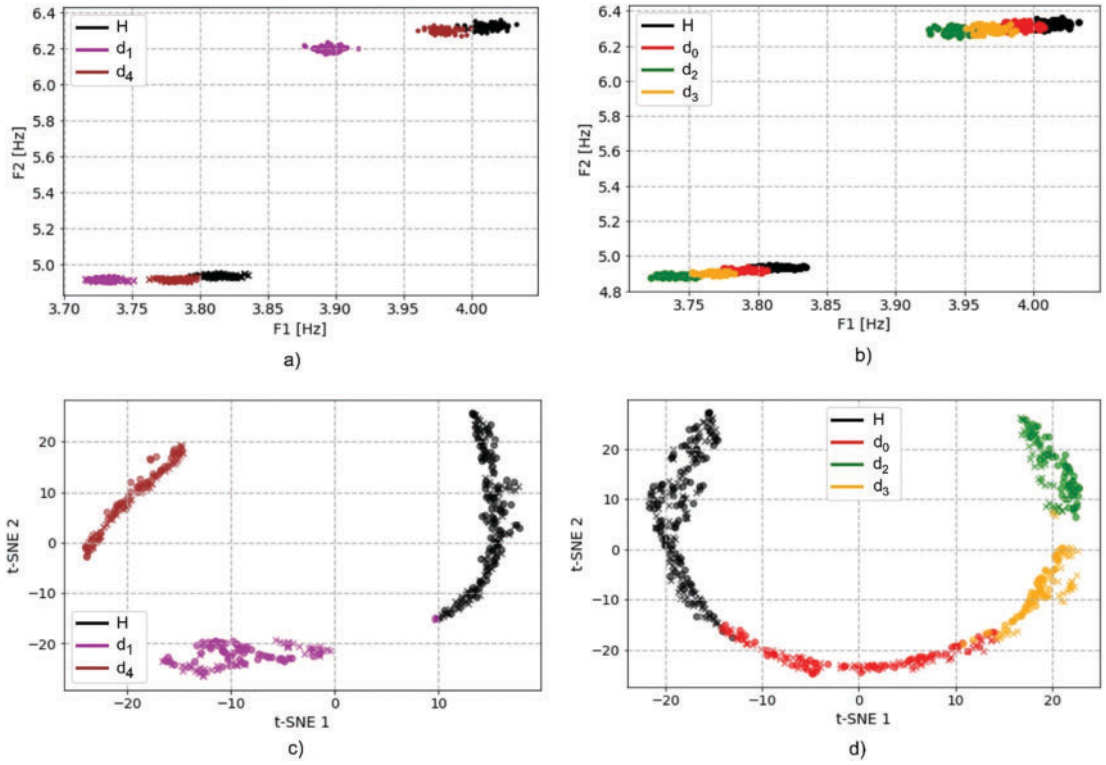

Figure 7. The original feature distribution of the Z24 FEM bridge (labelled with "×") and S101 FEM bridge (labeled with "◯") is shown in (a) and (b), referring to T2 and T3, respectively. Feature alignment post-DANN is shown for T2a (c) and T3a (d)

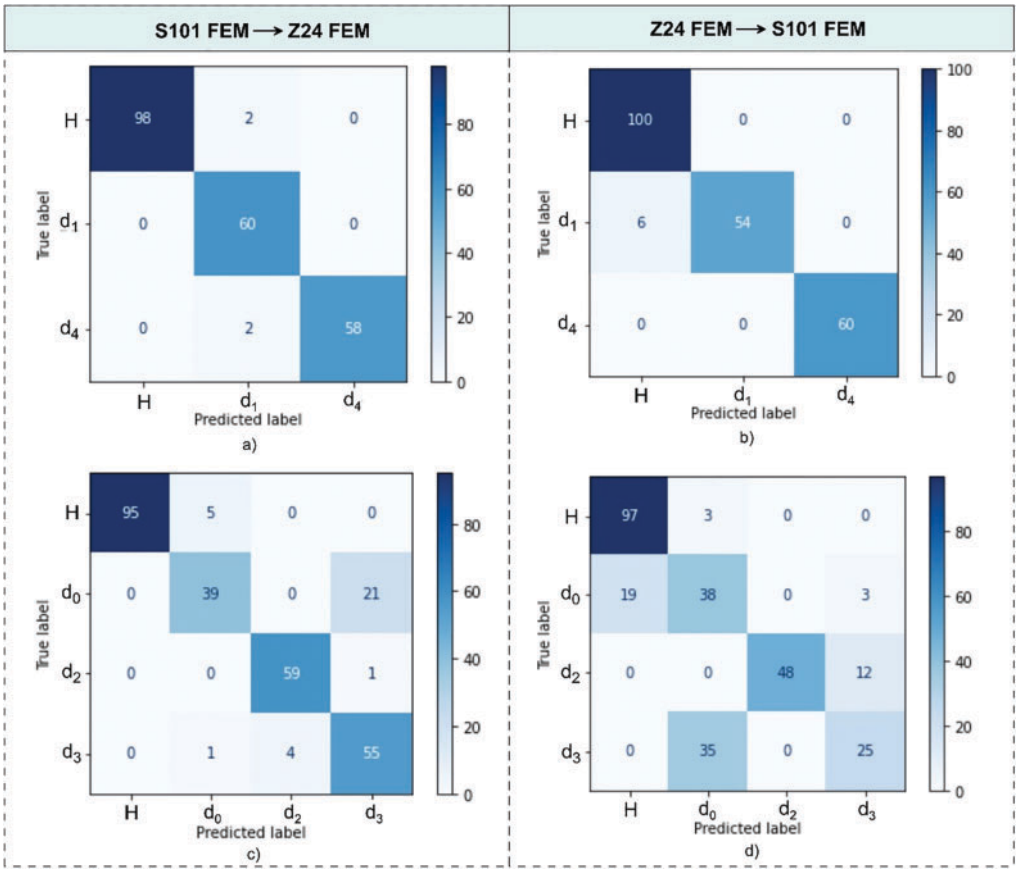

Figure 8. Damage classification performance via DANN for T2b (a), T2a (b), T3b (c), and T3a (d)

Table 4. Comparison of the performance, in terms of F1 score, between different DA techniques

Task	Source	NCA	JDA	DANN
T1	Z24	0.81	0.62	0.84
T2a	Z24	0.87	0.89	0.97
T2b	S101	0.82	0.93	0.98
T3a	Z24	0.30	0.69	0.71
T3b	S101	0.78	0.80	0.87

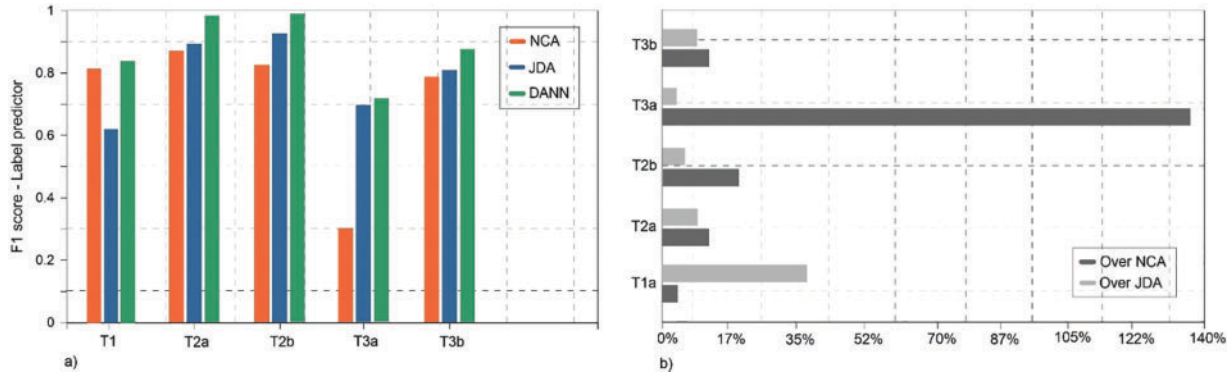

Figure 9. F1 score computed for each task, varying DA techniques (a), and the percentage increase in performance of the DANN framework compared to other methods (b)

plotted in Fig. 7d, referring in particular to T3a. However, for the T3 task, the algorithm's performance exhibited a decay, as can be inferred from the confusion matrices in Figs. 8c and 8d. Damage classes d_0 and d_3 proved particularly challenging to distinguish, likely due to the fact that they represent the same type of damage, differing only in severity. Despite the presence of false detection errors, the F1 score computed with the current methodology still improves over the other two techniques. To underline this statement, Table 4 summarizes the performance metrics for each task, considering different source domains and comparing the results obtained using (i) NCA alone, (ii) NCA combined with JDA, and (iii) the DANN-based framework. It is specified that NCA is used as a DA technique, i.e., case (i), and as a pre-processing technique to facilitate the application of JDA, i.e., case (ii), and DANN, i.e., case (iii).

In general, JDA provides a first improvement over NCA in most cases, with the exception of task T1, where JDA is not recommended for aligning the highly unbalanced datasets of the two real bridges.

Notably, the proposed network consistently outperforms the other methods across all tasks, achieving substantial percentage gains, as illustrated by the histograms in Fig. 9. Furthermore, DANN is proven to be more convenient (over JDA) in terms of computational costs, especially for large datasets such as the Z24 bridge. The importance of applying adversarial learning and performing DA is underlined by the poor results that are observed when training the classifier in the original frequencies space, yielding F1 scores between 0.30 and 0.50 for all the investigated tasks.

Conclusions

This study has demonstrated the potential of DANN for bridge health assessment, with a particular focus on TL-aided damage classification across two real bridges and two FEMs. By learning domain-invariant features, the proposed framework enables effective knowledge transfer, assuming the complete absence of labeled data in the target domain. The approach is especially well suited for large-scale monitoring scenarios involving multiple bridges, where collecting labeled data for each structure would be impractical.

In such contexts, DANN provides a scalable solution that enhances generalization and reduces reliance on manual labeling efforts. Beyond the specific context of bridge monitoring, TL inherently supports scalability, as it is designed to transfer knowledge from a well-labeled source domain (or multiple source domains) to one or more target domains with limited or unlabeled data. Any pair of comparable structures can potentially benefit from TL algorithms. However, defining the required level of similarity remains a challenge. A promising direction is to consider groups of structures with comparable typology, materials, or static schemes, and to systematically evaluate which characteristics most strongly influence TL performance.

In this work, natural frequencies associated with long-term continuous monitoring are selected as input features. However, to fully exploit the potential of DANN, future developments could consider using consecutive sequences of vibration data or combining heterogeneous input measurements. Since the nature of the input is closely tied to the architecture of the DANN, employing more advanced neural

network models, beyond simple MLPs, for each component may offer a valuable alternative, enabling better representation of complex structural behavior.

To simulate a variety of scenarios, several tasks are designed and addressed in this study, incorporating both real and simulated damage cases to explore different damage locations, severities, and extents. The results demonstrate high accuracy in damage identification and consistently improved performance across all tasks compared to previously applied DA methods, with the added benefit of reduced computational burden.

Working in the field of DA for SHM, one of the most critical aspects is the assumption of the same health-state scenarios across the network. Real-world situations can involve a very large number of possible damage classes, posing a significant challenge to classification accuracy and model scalability. Future work will need to address these aspects, potentially by integrating more efficient architectures and exploring open set DA, which represents a promising research direction to enhance robustness and adaptability in operational environments. Unlike standard DA techniques, open set DA explicitly accounts for the presence of unknown classes in the target domain, making it more suitable for realistic monitoring settings where new or unforeseen damage types may occur.

Overall, TL technologies, including the proposed DANN method, offer innovative strategies for identifying and classifying structural conditions based on knowledge learned from previously observed structures, potentially providing significant advantages for network-scale SHM applications.

Acknowledgments

This work was supported by the Italian Ministry of University and Research (MUR) through the funded project of national interest "TIMING—Time evolution laws for IMproving the structural reliability evaluation of existING post-tensioned concrete deck bridges" (Protocol No. P20223Y947).

References

[1] He Z, Li W, Salehi H, Zhang H, Zhou H, Jiao P. Integrated structural health monitoring in bridge engineering. *Autom Constr*. 2022;136(14):104168. doi:10.1016/j.autcon.2022.104168.

[2] Kamariotis A, Chatzi E, Straub D. A framework for quantifying the value of vibration-based structural health monitoring. *Mech Syst Signal Process*. 2023;184(5):109708. doi:10.1016/j.ymssp.2022.109708.

[3] Afshar A, Nouri G, Ghazvineh S, Hosseini L, Seyed H. Machine-learning applications in structural response prediction: a review. *Pract Period Struct Design Constr*. 2024;29(3):03124002. doi:10.1061/ppscfx.sceng-1292.

[4] Rashidi Nasab A, Elzarka H. Optimizing machine learning algorithms for improving prediction of bridge deck deterioration: a case study of Ohio bridges. *Buildings*. 2023;13(6):1517. doi:10.3390/buildings13061517.

[5] Wedel F, Marx S. Application of machine learning methods on real bridge monitoring data. *Eng Struct*. 2022;250(4):113365. doi:10.1016/j.engstruct.2021.113365.

[6] Weiss K, Khoshgoftaar TM, Wang D. A survey of transfer learning. *J Big Data*. 2016;3:1–40.

[7] Figueiredo E, Yano MO, Da Silva S, Moldovan I, Adrian Bud M. Transfer learning to enhance the damage detection performance in bridges when using numerical models. *J Bridge Eng*. 2023;28(1):28–21. doi:10.1061/(asce)be.1943-5592.0001979.

[8] Yano MO, Figueiredo E, Da Silva S, Cury A. Foundations and applicability of transfer learning for structural health monitoring of bridges. *Mech Syst Signal Process*. 2023;204(4):110766. doi:10.1016/j.ymssp.2023.110766.

[9] Poole J, Gardner P, Dervilis N, Bull L, Worden K. On statistic alignment for domain adaptation in structural health monitoring. *Struct Health Monitor*. 2023;22(3):1581–1600. doi:10.1177/14759217221110441.

[10] Gardner P, Bull LA, Gosliga J, Dervilis N, Worden K. Foundations of population-based SHM, Part III: heterogeneous populations-Mapping and transfer. *Mech Syst Signal Process*. 2021;149(9):107142. doi:10.1016/j.ymssp.2020.107142.

[11] Giglioni V, Poole J, Venanzi I, Ubertini F, Worden K. A domain adaptation approach to damage classification with an application to bridge monitoring. *Mech Syst Signal Process*. 2024;209(2):111135. doi:10.1016/j.ymssp.2024.111135.

[12] Giglioni V, Poole J, Mills R, Venanzi I, Ubertini F, Worden K. Transfer learning in bridge monitoring: laboratory study on domain adaptation for population-based SHM of multispan continuous girder bridges. *Mech Syst Signal Process*. 2025;224(3):112151. doi:10.1016/j.ymssp.2024.112151.

[13] Wu H, Li J, Zhang Q, Tao J, Meng Z. Intelligent fault diagnosis of rolling bearings under varying operating conditions based on domain-adversarial neural network and attention mechanism. *ISA Trans*. 2022;130(2):477–489. doi:10.1016/j.isatra.2022.04.026.

[14] Dai B, Frusque G, Li T, Li Q, Fink O. Smart filter aided domain adversarial neural network for fault diagnosis in noisy industrial scenarios. *Eng Appl Artif Intell*. 2023;126(9):107202. doi:10.1016/j.engappai.2023.107202.

[15] Jin Y, Song X, Yang Y, Hei X, Feng N, Yang X. An improved multi-channel and multi-scale domain adversarial neural network for fault diagnosis of the rolling bearing. *Control Eng Pract*. 2025;154(14):106120. doi:10.1016/j.conengprac.2024.106120.

[16] Martakis P, Reuland Y, Stavridis A, Chatzi E. Fusing damage-sensitive features and domain adaptation towards robust damage classification in real buildings. *Soil Dyn Earthq Eng*. 2023;166(1):107739. doi:10.1016/j.soildyn.2022.107739.

[17] Li ZD, He WY, Ren WX, Li YL, Li YF, Cheng HC. Damage detection of bridges subjected to moving load based on domain-adversarial neural network considering measurement and model error. *Eng Struct*. 2023;293(9):116601. doi:10.1016/j.engstruct.2023.116601.

[18] Li Z, Weng S, Xia Y, Yu H, Yan Y, Yin P. ross-domain damage identification based on conditional adversarial domain adaptation. *Eng Struct*. 2024;321(59):118928. doi:10.1016/j.engstruct.2024.118928.

[19] Maeck J, De Roeck G. Description of Z24 benchmark. *Mech Syst Signal Process*. 2003;17(1):127–131. doi:10.1006/mssp.2002.1548.

[20] Wenzel H, Veit-Egerer R, Widmann M. Case Study: S101. *Industr Saf Life Cycle Eng*. 2013:432–442.

[21] Döhler M, Hille F, Mevel L, Rücker W. Structural health monitoring with statistical methods during progressive damage test of S101 Bridge. *Eng Struct*. 2014;69(2082):183–193. doi:10.1016/j.engstruct.2014.03.010.
[22] Garcia-Macías E, Ubertini F. MOVA/MOSS: two integrated software solutions for comprehensive structural health monitoring of structures. *Mech Syst Signal Process*. 2020;143:106830. doi:10.1016/j.ymssp.2020.106830.
[23] Masciotta M-G, Ramos LF, Lourenco PB, Vasta M, De Roeck G. A spectrum-driven damage identification technique: Application and validation through the numerical simulation of the Z24 bridge. *Mech Syst Signal Process*. 2016;70(5):578–600. doi:10.1016/j.ymssp.2015.08.027.
[24] Eurocode 8. Eurocode 8—design of structures for seismic resistance—part 1: general rules, seismic actions and rules for buildings; 1998. https://eurocodes.jrc.ec.europa.eu/EN-Eurocodes/eurocode-8-design-structures-earthquake-resistance.
[25] Van der Maaten L, Hinton G. Visualizing data using t-SNE. *J Mach Learn Res*. 2008;9(11).