

From augmentation to translation: Data generation by conditional hierarchical variational autoencoder, enhancing monitoring mooring systems in floating offshore wind turbines

Original

From augmentation to translation: Data generation by conditional hierarchical variational autoencoder, enhancing monitoring mooring systems in floating offshore wind turbines / Fathnejat, Hamed; Nava, Vincenzo. - In: ENGINEERING APPLICATIONS OF ARTIFICIAL INTELLIGENCE. - ISSN 0952-1976. - 163:(2026). [10.1016/j.engappai.2025.112951]

Availability:

This version is available at: 11583/3004656 since: 2025-10-30T15:11:00Z

Publisher:

Elsevier

Published

DOI:10.1016/j.engappai.2025.112951

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



Research paper

From augmentation to translation: Data generation by conditional hierarchical variational autoencoder, enhancing monitoring mooring systems in floating offshore wind turbines

Hamed Fathnejat ^{a,b} ^{*,1}, Vincenzo Nava ^c 

^a Basque Center for Applied Mathematics, 14 Mazarredo Ave, Bilbao, 48009, Biscay, Spain

^b Construction Research Centre, National Research Council Canada, Building M-20, 1200 Montréal Road, Ottawa, K1A 0R6, Ontario, Canada

^c Politecnico di Torino, DIATI, Corso Duca degli Abruzzi, 24, Turin, Italy

ARTICLE INFO

Keywords:

Real-time structural health monitoring
Domain translation
Data augmentation
Conditional hierarchical variational autoencoder
Diffusion process
Floating offshore wind

ABSTRACT

The integrity of mooring systems in floating offshore wind turbines (FOWTs) is crucial, as their degradation alters the platform's dynamic behavior. A robust machine learning-based health monitoring system that continuously monitors different mooring systems for FOWTs requires data under diverse health, operational, and metocean conditions. To this end, we propose a Conditional Hierarchical Variational Autoencoder (CHVAE) generative model designed for simultaneous data augmentation and domain translation to generate the required data. We train the model to learn the nonlinear relationships between healthy and minority-damaged fairlead tension records from the source mooring system across various sea states. CHVAE generates realistic damaged responses under diverse conditions by leveraging healthy data from the target mooring system. We first assess CHVAE's ability to augment minority data based on majority distribution, validated on the Modified National Institute of Standards and Technology (MNIST) benchmark dataset. This experiment compares the performance of CHVAE variants with conventional and recent oversampling methods. Second, the open-source software OpenFast simulates the testing and training datasets for simultaneously data augmentation and domain translation on the Offshore Code Comparison Collaboration Continuation (OC4) semi-submersible platform (DeepCwind) FOWT benchmark. OpenFast and CHVAE records are compared through visual, statistical, and behavioral methodologies. Simulations utilize diverse wave seeds to represent excitation randomness and undetected damage severities, assessing CHVAE's one-to-all capability. Generated records for unobserved sea states and damage severities closely mimic real behavior in downstream binary classification, illustrating the versatility of CHVAE for zero-shot, real-time damage identification.

1. Introduction

Wind energy is essential for achieving zero-emissions electricity, and floating offshore wind turbines (FOWTs) have considerable promise for capturing additional renewable energy in deep oceans (Williams et al., 2024). Fig. 1 depicts a typical FOWT configuration, highlighting its floating platform and mooring system, enabling operation in deeper waters compared to conventional fixed-bottom turbines.

Until 2022, 32 FOWTs have been installed at a commercial scale worldwide (Williams et al., 2024) for a total capacity of 121 MW (Edwards et al., 2023). Nevertheless, the outlook for the next decade is

that several projects are expected to be successfully deployed, reaching a capacity of 18.9 GW by 2030 and 264 GW by 2050 (DNV, 2022). Among the diverse concepts of floating support structures, semisubmersible platforms can represent a valuable solution when stabilization needs to be achieved by the waterplane (Chen and Kim, 2021; Edwards et al., 2023). Mooring systems are critically important in floating semisubmersible platforms because they provide station-keeping and operational functionality ((Chen and Kim, 2021). Further, they also considerably impact the project's economic viability (Nava et al., 2019).

* Corresponding author.

E-mail address: Hamed.Fathnejat@nrc-cnrc.gc.ca (H. Fathnejat).

¹ Present address: Construction Research Centre, National Research Council Canada (NRC), 1200 Montreal Road, Ottawa, ON K1A 0R6, Canada.




Fig. 1. Typical configuration of a floating offshore wind turbine (FOWT), depicting the floating platform, the turbine, and mooring system that facilitates operation in deep-water locations.

Although FOWTs hold great promise, a major challenge lies in accurately simulating their behavior and monitoring the conditions of their mooring systems, particularly in developing a digital twin (DT) for these systems (Longman et al., 2023). The limited number of installed platforms makes it difficult to gather data, especially for damaged conditions. Additionally, simulating mooring system behavior using physical models is both computationally intensive and prone to uncertainty. Addressing these challenges requires innovative methods to fully understand and manage FOWT mooring systems.

In structural health monitoring (SHM) of mooring systems, data on various environmental and operational conditions are often unbalanced and insufficiently organized (Tamuly et al., 2025; Fathnejat and Nava, 2025). Data from damaged states are particularly scarce and less systematically collected than those from healthy conditions, presenting a major challenge when developing condition monitoring systems (Sharma and Nava, 2024; Fathnejat et al., 2023; Altabay et al., 2023; Omid and Liang, 2019; Rezaniaeie Aqdam et al., 2018; Ghiasi et al., 2019; Fathnejat et al., 2014; Ahmadi-Nedushan and Fathnejat, 2022; Fathnejat and Ahmadi-Nedushan, 2017; Torkzadeh et al., 2016; Fathnejat and Ahmadi-Nedushan, 2020). This challenge is compounded by the common scenario in which a new mooring system has just been deployed and is still in a healthy state, making it unrealistic to expect comprehensive damaged data that spans the full range of operational sea states, such as waves, currents, and wind. Even when damaged data are available, variability in sea states can further complicate accurate detection, potentially degrading performance. Addressing these issues is critical for advancing the reliability and effectiveness of SHM in floating offshore wind turbines.

Recent advancements in artificial intelligence (AI) have led to its growing application in marine and offshore energy systems, demonstrating considerable potential for monitoring, control, and diagnostic functions in uncertain contexts and restricted data availability. Machine learning has been applied in wave energy applications to decrease operational costs and improve diagnostic accuracy, utilizing real experimental data from facilities like the Mutriku Wave Power Plant (M'zoughi et al., 2024b). Metaheuristic control strategies have been investigated

to mitigate platform vibrations in hybrid systems that integrate oscillating water columns and floating wind turbines (M'zoughi et al., 2024a). AI-based methods have been established for floating wind turbines, focusing on real-time prediction of turbine thrust (Jiang et al., 2019), forecasting of motion and mooring loads using neural networks (Medina-Manuel et al., 2024), and the development of AI-in-the-loop testing platforms to support digital twin-based validation (Jiang et al., 2024). The contributions underscore the increasing significance of AI in offshore system analysis and establish a basis for further study into intelligent SHM solutions in data-scarce conditions.

Domain adaptation (DA) methods have been proposed to address this challenge. These methods transfer knowledge from a source domain with sufficient labeled data to a target domain with limited or no labeled data, even if the data distributions are different but related. Recent DA techniques have been applied for damage detection in structural health monitoring (Ghiasi et al., 2025; Chen et al., 2023; Gigliani et al., 2024; Ragab et al., 2021). However, as noted in Ragab et al. (2021), a significant challenge in using DA is the noticeable decline in detection performance, mainly when shallow neural networks are employed as classifiers. Despite progress in DA methods, the accuracy of damage detection can still decrease significantly, sometimes dropping to 50% or lower. This decline in accuracy highlights the need for further research into more robust DA techniques to ensure high detection accuracy, even in challenging conditions where target domain data is scarce and may differ from the source domain.

To overcome the drop in accuracy, researchers seek to utilize the deep generative models (DGMs) properties for DA and domain translation. One framework for implementing this utilization involves training DGMs to enhance damage-associated data under various operational and environmental conditions using minority data from a source domain (data augmentation), followed by applying the trained DGMs in the source domain to estimate the features of corresponding damage-associated data in the target domain (domain translation). Fig. 2 illustrates this framework for data augmentation and domain translation, whereby the model gains the ability to generate minority data in

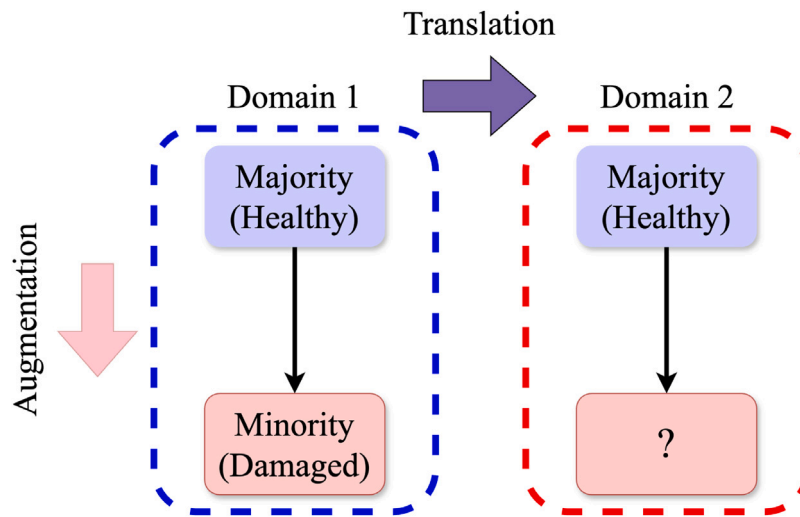


Fig. 2. Utilizing majority data for cross-domain minority data generation.

domain 2 by utilizing the knowledge acquired during training for data augmentation in domain 1.

In this context, Kiranyaz et al. (2024) offers a zero-shot bearing fault detection approach capable of identifying real faults in a new machine, irrespective of its operational conditions, sensory configuration, or fault intensities. The approach employs 1D Operational Generative Adversarial Networks (Op-GANs) to delineate the inherent transition from a healthy to a faulty state of bearing within the signal domain.

Data augmentation can function as a domain adaptation approach when generating a broader range of various environmental and operational damage-associated data. Data augmentation is one of the potential approaches that can be adopted to solve the issue of imbalanced data distributions. Data distribution imbalance can be corrected through over-sampling of minority (damaged) data or under-sampling of majority (healthy) data. The application of DGMs represents an innovative approach for resampling, understanding the distribution and behavior of minority data, and augmenting it with synthetic data. Studies have shown that in addressing imbalanced datasets, conditional VAEs (CVAEs) surpass conditional GANs (CGANs) (Fajardo et al., 2021); the FID analysis in this research indicates that the CVAEs presented are more resilient to unbalanced datasets compared to their CGAN counterparts, a finding that is inverted when GANs and VAEs are applied to the original, well-balanced MNIST and Fashion MNIST datasets. Currently, VAEs are mostly utilized to address several issues within the SHM domain, including feature extraction and data dimensionality reduction (Coraça et al., 2023; Ma et al., 2020; Pollastro et al., 2022; Mylonas et al., 2020; Lee et al., 2023).

In the context of class-imbalanced SHM problems, integrating a focus-loss-optimized CNN classifier with VAE-GAN-based augmentation, a new approach called VGAIC-FDM improves fault identification in the presence of unbalanced data. The result is a high level of diagnostic accuracy and F1-scores (Li et al., 2024). Zhao et al. (2022) developed the normalized conditional VAE with adaptive focal loss (NCVAE-AFL) framework to improve diagnostic efficacy for minority classes in class-imbalanced structural health monitoring challenges. Li and Betti (2023) introduced an innovative data augmentation system utilizing a conditional VAE architecture to deliver cepstral coefficients as response characteristics.

The papers mentioned earlier solely present techniques aimed at mitigating the class imbalance problem in deep learning-based damage identification and classification. However, VAE has not yet been applied to generate real-scale time-series minority data. Additionally, DGMs are trained solely on the distribution of minority data, using only information from the minority class. A recent study indicates that

DGMs can enhance minority data by increasing diversity and richness while employing majority data sets. In this context, Ai et al. (2023) presented an objective function that allows VAEs to enhance minority data by utilizing majority-based priors. We propose a novel conditional hierarchical variational approximation to parameterize a diffusion process, aimed at constructing the minority data distribution by leveraging features from both majority and minority data. The hierarchical VAE (HVAE) improves the VAE architecture through the incorporation of multiple stochastic latent layers or distributions (Havtorn et al., 2021). The diffusion process represents a stochastic mechanism that alters data progressively over time (Sohl-Dickstein et al., 2015). Our proposed method employs a two-stage pretrain-finetune training framework to generate and augment real-scale damaged state time series (minority data) from a healthy state (majority data) across various environmental and operational conditions.

To verify our proposed DGM model, we used the MNIST benchmark dataset (Lecun et al., 1998), comparing results to those in Ai et al. (2023), where conventional and recent oversampling methods are evaluated on MNIST data augmentation. Additionally, we maintained the same conditions and multiclass classification model as Ai et al. (2023) to ensure a fair comparison.

Due to the limited availability of real-world data for mooring systems for FOWTs, the datasets analyzed in this paper are generated using OpenFast (OpenFAST, 2023), based on the OC4 platform's numerical model developed under the DeepCwind project (Robertson et al., 2014a). This model was selected as the source system, as it has been validated against experimental data and Computational Fluid Dynamics simulations (see for example, Gorostidi et al., 2023). We chose an alternative mooring configuration for the target system providing an equivalent pretension, but characterized by different mass per unit length and axial stiffness calculated as described in the Orcaflex manual for studless chains (Orcina, 2022).

The sections of our study are organized as follows: Section 2 outlines the methodology, including the proposed framework and its theoretical foundations. Section 3 details the experiments conducted on CHVAE, covering implementation specifics, evaluation criteria, and outcomes analysis. This section highlights the effectiveness of the proposed framework in various contexts of data augmentation and domain translation. Finally, Section 4 summarizes the findings, concluding the study.

2. Methodology

This research proposes a novel framework that utilizes the principles of conditional variational autoencoder (CVAE) algorithms and diffusion

probability models (DPM) to reconstruct the latent distribution of minority data (data augmentation) and essentially facilitates the creation of a real-time monitoring system utilizing domain translation.

2.1. Pre-trained VAE with two decoders

When reconstructing real-scale data, especially in the case of time-series data, the VAE requires access to denormalization parameters, as normalized data is essential for optimal training. This challenge becomes more prominent in domain translation, where there is no information about the damage-associated target mooring data and their denormalization parameters. To tackle this challenge, we have introduced a second decoder, which is conditioned on denormalization parameters; moreover, to sample from the latent space to estimate the denormalization parameters by this decoder, building the latent space in the encoder is also conditioned on denormalization parameters. Hence, during the training, the first decoder learns to reconstruct the normalized input data, and simultaneously, the second conditional decoder learns to reconstruct and estimate corresponding denormalization parameters from conditional latent space.

pre-training is conducted on the majority class, which in this case consists of healthy data from the source mooring system. The trainable weights of the proposed VAE are optimized based on the Eq. (1):

$$\text{ELBO}_{\theta, \theta', \phi}(X^+, d_1) = \mathbb{E}_{q_\phi} [\log p_\theta(X^+ | z_1)] + \mathbb{E}_{q_\phi} [\log p_{\theta'}(d_1 | z_1, d_1)] + \beta \times (-\text{KL} [q_\phi(z_1 | X^+, d_1) \parallel p_\theta(z_1 | d_1)]) \quad (1)$$

In Eq. (1), X^+ is normalized healthy data of source mooring and d_1 is the denormalization matrix corresponding to X^+ . β serves as a hyper-parameter that governs the equilibrium between reconstruction loss and KL divergence during the training process of the VAE.

2.2. Fine-tuned VAE with two conditional decoders

To control the data generation corresponding to different damage severities, a conditional fine-tuned VAE is designed and trained on source mooring data to generate damage-associated data of target mooring with various severities. In this regard, two conditional decoders are designed for the fine-tuned VAE. The first one is conditioned on damage severity/state (S) to learn how to generate corresponding minority damage-associated normalized data and the second one is conditioned on denormalization parameters of healthy data (d_1) from the source mooring system. In the case of domain translation, since we lack information from damage-associated data of the target mooring, we have designed the second decoder to learn and estimate the denormalization parameter (d_2) based on the denormalization parameter of healthy data (majority class), (d_1). Then, inspired by the diffusion process, the conditional hierarchical VAE (CHVAE) is assumed to have two stochastic layers constructing two latent variables z_1 and z_2 . Through the introduction of a variational approximation to the true posterior and bottom-up formula for the inference model, by extending the loss function of the standard VAE, we can derive Eq. (2):

$$\begin{aligned} \text{ELBO}_{\theta, \phi}(x^-) &= \mathbb{E}_{q_\phi} [\log p_\theta(x^-, z_2, z_1) - \log q_\phi(z_2, z_1 | x^-)] \\ \text{ELBO}_{\theta, \phi}(x^-) &= \mathbb{E}_{q_\phi} [\log p_\theta(x^- | z_2) + \log p_\theta(z_2 | z_1) + \log p_\theta(z_1) \\ &\quad - \log q_\phi(z_2 | x^-) - \log q_\phi(z_1 | z_2)] \end{aligned} \quad (2)$$

In Eq. (2), x^- represents the minority-damaged state data, while z_1 denotes the latent distribution corresponding to the healthy or baseline condition, and z_2 represents the latent distribution associated with the damaged state.

A DPM is a specific type of latent variable model that represents data using a sequence of \mathcal{T} variables (Ho et al., 2020). The initial variable $x_0 \sim p(x)$ corresponds to the observed data, while the remaining variables $x_{1:\mathcal{T}}$ are hidden and not directly observed.

The DPM model defines a joint distribution $p_\theta(x_{0:\mathcal{T}})$ that reverses a diffusion process. This process is a Markov chain with Gaussian transitions that start from $p(x_\mathcal{T}) = \mathcal{N}(x_\mathcal{T}; 0, I)$. The model learns how to reverse the process by optimizing a variational bound, as defined in Eq. (3):

$$-\text{ELBO}_{\theta, \phi}(x_0) = \mathbb{E}_{q_\phi} \left[-\log p_\theta(x_\mathcal{T}) - \sum_{t \geq 1} \log \frac{p_\theta(x_{t-1} | x_t)}{q_\phi(x_t | x_{t-1})} \right] \quad (3)$$

For $\mathcal{T} = 2$ in Eq. (3), $x_\mathcal{T}$ and x_0 are equivalent to z_1 and x^- , respectively; thus, based on Eqs. (2) and (3), the proposed loss function can be presented as Eq. (4):

$$\begin{aligned} \text{ELBO}_{\theta, \phi}(x^-) &= \mathbb{E}_{q_\phi} [\log p_\theta(x^- | z_2)] - \text{KL} [q_\phi(z_2 | x^-) \parallel p_\theta(z_2 | z_1)] \\ &\quad - \text{KL} [q_\phi(z_1 | z_2) \parallel p_\theta(z_1)] \end{aligned} \quad (4)$$

By adding the second conditional decoder and considering the conditional fine-tuned VAE, we can derive Eq. (5):

$$\begin{aligned} \text{ELBO}_{\theta, \theta', \phi}(x^-, d_2, S) &= \mathbb{E}_{q_\phi} [\log p_\theta(x^- | z_2, S)] + \mathbb{E}_{q_\phi} [\log p_{\theta'}(d_2 | z_2, d_1, S)] \\ &\quad - \text{KL} [q_\phi(z_2 | x^-, d_2, S) \parallel p_\theta(z_2 | z_1, d_2, S)] \\ &\quad - \text{KL} [q_\phi(z_1 | z_2, d_1) \parallel p_\theta(z_1 | d_1)] \end{aligned} \quad (5)$$

Based on the fine-tuning method proposed in Eq. (5), to obtain z_2 , the distributions $q_\phi(z_1 | z_2, d_1)$ and $q_\phi(z_2 | x^-, d_2, S)$ are simultaneously estimated. This is done to ensure that these distributions approximate $p_\theta(z_1 | d_1)$ and $p_\theta(z_2 | z_1, d_2, S)$ as closely as possible, respectively. Additionally, when z_2 and S are decoded, it should allow for the reconstruction of x^- , i.e., $p_\theta(x^- | z_2, S)$. Since z_1 and z_2 share the same dimension (D), $q_\phi(z_1 | z_2, d_1)$ follows a standard normal distribution based on the assumptions used in the reparameterization of the pre-trained VAE (Eq. (1)). Consequently, the second KL divergence term in Eq. (5) remains constant throughout the fine-tuning process. Hence, we introduced a novel reparameterization for fine-tuning the CHVAE as stated in Eq. (6); and, the trainable weights of the CHVAE are optimized through the minimization of the loss function presented in Eq. (7):

$$z_2 = \left[\mu_1 + e^{\frac{\log(\sigma_1^2)}{2}} \right] \times \left[\mu_2 + e^{\frac{\log(\sigma_2^2)}{2}} \right] \times z_1 \quad (6)$$

In Eq. (6), μ_1 and μ_2 are the means and, $\log(\sigma_1^2)$ and $\log(\sigma_2^2)$ are the log variances of normal distributions of z_1 and z_2 , respectively.

$$\begin{aligned} \mathcal{L}_{\text{CHVAE}}(x^-, d_2, S; \phi, \theta, \theta') &= \mathcal{L}_{\text{recon}, \theta}(x^-, \hat{x}^-) + \mathcal{L}_{\text{recon}, \theta'}(d_2, \hat{d}_2) \\ &\quad + \beta \times (-\text{KL} [q_\phi(z_2 | x^-, d_2, S) \parallel p_\theta(z_2 | z_1, d_2, S)]) \\ &\quad + \text{FID}_{\text{loss}} \end{aligned} \quad (7)$$

The term $\mathcal{L}_{\text{recon}}$ in the first and second terms of Eq. (7) quantifies the differences between the input x^- , d_2 and their reconstructions \hat{x}^- , \hat{d}_2 , respectively. The reconstruction error function in a VAE is determined by the properties of the data and the activation function utilized in the decoder's last layer. Binary Cross-Entropy (BCE) is frequently employed for binary or binarized data, especially in conjunction with a sigmoid activation function. Conversely, mean squared error (MSE) is better for continuous data, such as time-series data, generally utilizing tanh or linear activation functions. This option improves model performance by synchronizing reconstruction loss with data characteristics and use tasks.

The MSE is defined by Eq. (8):

$$\text{MSE}(x^-, \hat{x}^-) = \frac{1}{N} \sum_{i=1}^N (x_i^- - \hat{x}_i^-)^2 \quad (8)$$

In Eq. (8), x_i^- are the real input minority data, \hat{x}_i^- are the reconstructed data, and N is the number of data points.

The Binary Cross-Entropy (BCE) is defined by Eq. (9):

$$\text{BCE}(x^-, \hat{x}^-) = -\frac{1}{N} \sum_{i=1}^N [x_i^- \log(\hat{x}_i^-) + (1 - x_i^-) \log(1 - \hat{x}_i^-)] \quad (9)$$

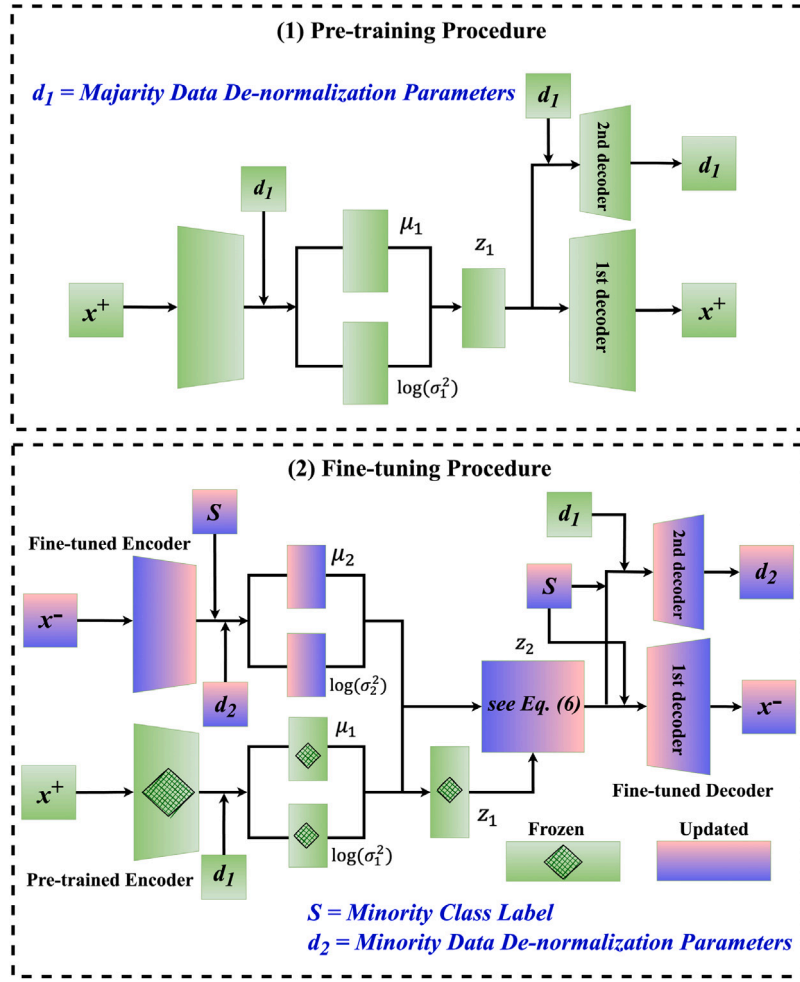


Fig. 3. CHVAE training framework.

In Eq. (9), x_i^- and \hat{x}_i^- represent the true binary values and predicted probabilities, respectively, and N is the number of data points.

Furthermore, the Fréchet Inception Distance (FID) (Heusel et al., 2017) loss is calculated via Eq. (10):

$$\text{FID}_{\text{loss}} = \|\mu_{x^-} - \mu_{\hat{x}^-}\|_2^2 + \text{Tr}\left(C_{x^-} + C_{\hat{x}^-} - 2(C_{x^-}C_{\hat{x}^-})^{1/2}\right) \quad (10)$$

In Eq. (10), μ_{x^-} and $\mu_{\hat{x}^-}$ are the means and C_{x^-} and $C_{\hat{x}^-}$ are the covariance matrices of the real and generated minority data, respectively. Here, Tr denotes the trace of the matrices. The FID is used as an index to monitor the fine-tuning process, with lower FID values indicating more similarity between the compared data. The training framework of the proposed DGM, CHVAE, is described in Fig. 3.

3. Experiments

In this section, we first evaluate the data augmentation quality of CHVAE by applying it to a multiclass classification task using the MNIST benchmark image dataset (Lecun et al., 1998) as a downstream application. Additionally, we compare the performance of the proposed framework with MGVAE, introduced by Ai et al. (2023). Second, we apply CHVAE to construct a real-time monitoring system based on domain translation for various mooring systems of floating offshore wind turbines (FOWTs).

This evaluation commences by assessing the proposed CHVAE's performance and generalizability across two distinct data modalities.

The MNIST dataset, composed of image data, is a recognized benchmark in generative modeling, facilitating standardized assessment and

comparison with contemporary deep generative models such as MGVAE across different class imbalance scenarios.

The FOWT dataset consists of sequential time-series data from mooring line tensions, illustrating a complex real-world scenario in SHM.

This two-stage experimental design enables the initial assessment of CHVAE's capacity to learn compact and expressive latent representations in a controlled environment, followed by its applicability in a complex domain translation task pertinent to FOWT monitoring. CHVAE's robust performance in image and time-series modalities demonstrates its versatility and potential for broader application in structural health monitoring contexts.

3.1. MNIST data augmentation

We assumed digits 0 through 4 as the majority classes, each containing 6000 samples, resulting in 30,000 samples for the majority. The minority classes digit 5 through 9, were downsampled to 300 samples (60 per class) and 50 samples (10 per class), corresponding to imbalance ratios $\rho = 100$, and $\rho = 600$, respectively, the same as (Ai et al., 2023). The downstream classifier is a fully connected neural network with two layers, the same as one used in Ai et al. (2023).

3.1.1. CHVAE specifications

For a thoroughly fair comparison, the architecture of the VAEs employed for pre-training and fine-tuning is identical to that of the MGVAE in Ai et al. (2023), which utilizes an MLP-based architecture.

To improve the comparative experiment, we examined the hierarchical VAE (HVAE) without conditioning on minority classes, S in

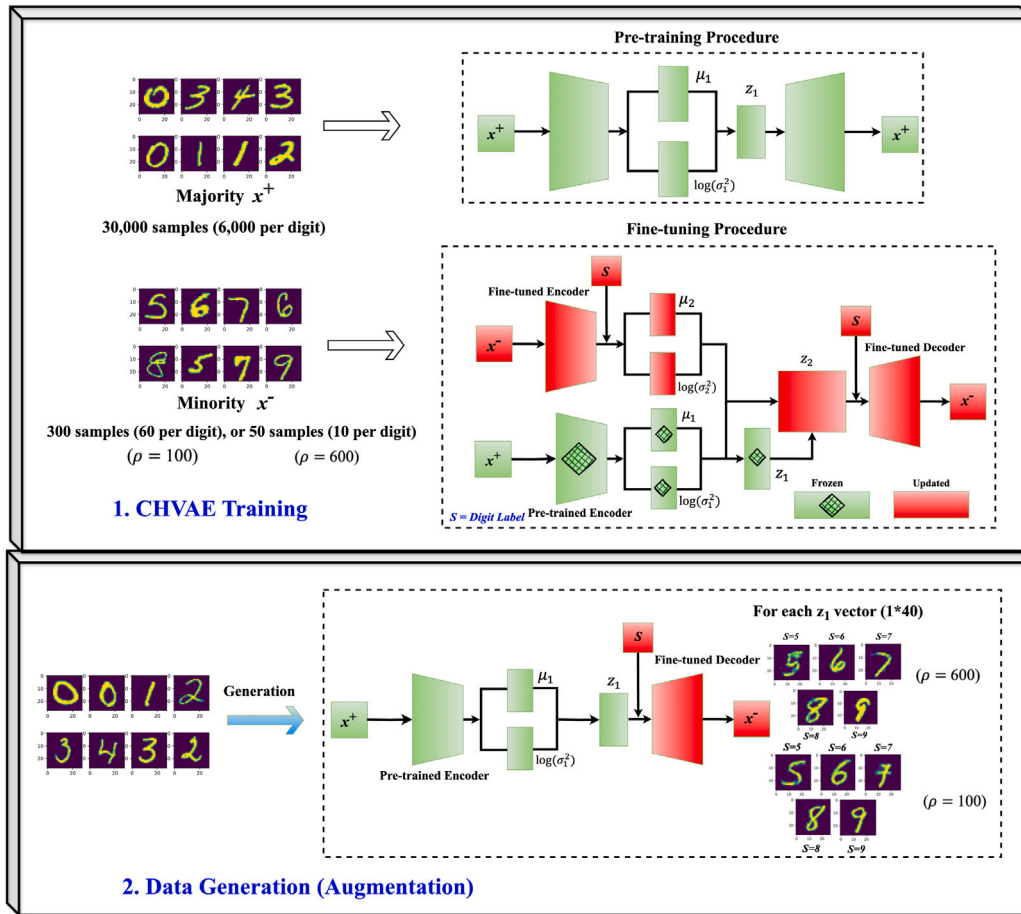


Fig. 4. The visualization of CHVAE framework for MNIST data augmentation.

Eq. (5), during the fine-tuning process. Therefore, Eq. (4) is the loss function in this case. To accomplish this, similar to the most recent method MGVAE (Ai et al., 2023), which has demonstrated superior outcomes to date, we implement a fine-tuning procedure for each minority class. Consequently, we develop a separate fine-tuned HVAE for each minority class. In CHVAE, we implement a distinct fine-tuning procedure to fine-tune a single model for all minority classes.

The pre-trained and fine-tuned MLP-based VAE architecture, identical to that in Ai et al. (2023), utilized in HVAE and CHVAE training is as follows:

Encoder Architecture:

Input : (Batch size, 28×28)

DenseLayer : 300,
activation = LeakyReLU($\alpha = 0.2$)

DenseLayer : 300,
activation = LeakyReLU($\alpha = 0.2$)

BatchNormalization
 μ = Dense: Latent dim. = 40,
 σ = Dense: Latent dim. = 40,
activation = None

Decoder Architecture:

Input : (Batch size, 40)

DenseLayer : 300,
activation = ReLU

BatchNormalization

DenseLayer : 300,

activation = ReLU

BatchNormalization

Output : (Batch size, 28×28),

activation = Sigmoid

The reconstruction loss function used is Binary Cross-Entropy (BCE) (Eq. (9)), and the optimizer applied is Adam, set with a learning rate of 0.001. The pre-training batch size is set to 32 for 100 epochs. The batch size assigned for fine-tuning each minority class in HVAE is as follows: 60 ($\rho = 100$) over 10,000 epochs and 10 ($\rho = 600$) for 10,000 epochs. For CHVAE, we need to fine-tune a single CHVAE for all minority classes while incorporating a conditioning layer in both the encoder and decoder, as illustrated in Fig. 4. The fine-tuning epochs for $\rho = 100$ are set to 100,000 with a batch size of 300, whereas for $\rho = 600$, it is 150,000 with a batch size of 50.

Furthermore, the cyclical annealing schedule, as established by Fu et al. (2019), was employed in MLP-based CHVAE to determine the optimal β values in Eq. (7), thereby enhancing the training quality during the fine-tuning stage. We establish the minimum value of β at 0 and the maximum value at 2×10^{-4} for applying this technique. The β value is reset every 20 epochs, according to the computed cycle length. The fine-tuning of the MLP-based HVAE, the CNN-based CHVAE, and the pre-training VAE for all HVAE and CHVAE models utilizes a constant β value of 5×10^{-4} , as specified in Eq. (1).

To further enhance the outcomes, we additionally implemented CNN-based VAEs for CHVAE. The visualization of our proposed framework is presented in Fig. 4.

Table 1
The CNN-based VAE parameters (pre-training and fine-tuning procedures) for MNIST data.

Parameter	Value
Batch size	pre-train:32, fine-tune: ($\rho = 100$)=300, ($\rho = 600$)=50
Padding	Same
Activation function (Output layer of decoder)	Sigmoid
Optimizer	Adam
Learning rate	0.001
Reconstruction loss function	BCE (Eq. (9))
Latent space dimension	40

Based on Fig. 4, the pre-trained CNN-based VAE architecture, which is trained on the majority (x^+), consists of an encoder with five 2D CNN layers with (3×3) kernel size, having 32, 64, 128, 256, and 512 filters, respectively with an input shape of (Batch size, 28, 28, 1). The stride is 1, except for the last CNN layer where it is 2. The activation function applicable to all layers, except output, is ReLU. MLP layers generate latent variables. For the decoder, five Conv2DTranspose layers are used, mirroring the features of the encoder.

The fine-tuned CNN-based VAE architecture is as follows:

Encoder: Input01(Batch size, 28, 28, 1), Conv2D(32, (3, 3), stride(1, 1)), Conv2D(64, (3, 3), stride(2, 2)), Conv2D(64, (3, 3), stride(1, 1)), Conv2D(128, (3, 3), stride(1, 1)), Flatten, MLP(64), Input02(S), Activation function applicable to all layers: ReLU;

Decoder: Input01(z_2), Input02(S), MLP($14 \times 14 \times 128$), Reshape(14, 14, 128), Cov2DTranspose(128, (3, 3), stride(1, 1)), Cov2DTranspose(64, (3, 3), stride(1, 1)), Cov2DTranspose(64, (3, 3), stride(1, 1)), Cov2DTranspose(32, (3, 3), stride(2, 2)).

The activation function applies to all layers except the output: ReLU. Other VAE characteristics are mentioned in Table 1.

3.1.2. Qualitative and quantitative assessment

The qualitative comparison of minority data generation results using MLP-based HVAE, MLP-based CHVAE, and CNN-based CHVAE is illustrated in Fig. 5.

Fig. 5 presents a comparative visualization of majority real data (digits 0–4, top row) and minority generated data (digits 5–9, bottom row) for MNIST at imbalance ratios of $\rho = 100$ and $\rho = 600$. Moreover, Fig. 5 illustrates the stylistic similarities between the generated minority data and the majority real data regarding features such as tilt angle, thickness, and overall structure. We provide both the reference majority sample and the associated minority sample, as the HVAE generation procedure is one-to-one. Alternatively, CHVAE is capable of generating one-to-all.

We can observe that:

- **Stylistic Consistency:** The generated minority digits closely resemble the majority data, maintaining stylistic coherence regarding tilt and thickness. This reflects the ability of all HVAE models, regardless of the presence of conditional feature, to capture essential data characteristics, ensuring realistic augmentation.
- **Generation Quality Across ρ :**
 - At $\rho = 100$, the generated digits generated by CNN-based CHVAE are highly consistent with majority data in terms of tilt angle, thickness, and overall quality, yielding the best downstream multiclass classification results, suggesting that the model effectively learns diverse and representative patterns for minority class augmentation (Table 3).
 - At $\rho = 600$, the MLP-based HVAE generated digits maintain acceptable quality and stylistic consistency, and exhibit the best performance in downstream multiclass classification task as presented in Table 3, surpassing both MLP-based CHVAE and CNN-based CHVAE.

Table 2
The quality evaluation of generating minority MNIST data by CHVAE.

Imbalance ratio	FID	MSE
$\rho = 100$	0.0241	0.0978
$\rho = 600$	0.0498	0.0946

The unconditional variant of HVAE could provide superior results by fine-tuning each HVAE specifically for each minority class. On the other hand, designing a conditional variant of HVAE presents challenges, as it requires additional epochs and a more complex architecture (CNN-based) to fine-tune a single CHVAE for generating various minority classes by controlling the condition as a minority class label. However, the conditional type of HVAE has significance in SHM, particularly for generating a specific minority class (associated with particular damage severity) from a designated majority class (healthy).

To emphasize the practical benefits of employing CHVAE for generating minority class data, the qualitative outcomes of MNIST data augmentation with CNN-based CHVAE are illustrated in further detail in Fig. 6. The first column in Fig. 6(a) and (b) displays the reference majority sample, while the subsequent columns present the corresponding generated minority samples based on S value.

As illustrated in Fig. 6 and also Fig. 4, the data generation section, for each z_1 value (e.g., using digit 1 from the majority class, which includes digits 0, 1, 2, 3, and 4), we can generate all digits in the minority dataset (5, 6, 7, 8, and 9). By altering the input digit, different minority digits (5 through 9) can be generated by defining S .

As shown in Fig. 6, the key advantage of the proposed model over MGVAE is that we can generate all the minority classes by fine-tuning the CHVAE model just once. As a result, we can generate the corresponding image for any minority class simply by selecting the class label S in Eq. (5) and using the same location (z_1 value) in latent space. Generating diverse minority samples from each majority sample enhances balance and representation within the dataset.

Furthermore, as illustrated in Fig. 6(c), CHVAE can generate inter-class minority data that can be utilized for generating damage-associated records with varying levels of damage severity.

Based on Eq. (4), the architecture of hierarchical VAE can be sketched as Fig. 7 and based on this architecture, after fine-tuning, the latent distribution of majority data, z_1 , is a baseline distribution to generate diverse minority data in our proposed framework. In Fig. 7 we visualize the z_1 and z_2 , which are constructed by feeding majority and minority data to the encoder of pre-trained VAE and the encoder of fine-tuned CHVAE, respectively, as illustrated in Fig. 4, utilizing the t-distributed stochastic neighbor embedding (t-SNE) (Van Der Maaten and Hinton, 2008). T-SNE is applied to embed 40-dimension z_1 and z_2 to 3-dimension. The designed conditional fine-tuned decoder can easily sample from z_1 by defining arbitrary S values to generate diverse minority data (x^-).

The statistical index, FID, based on Eq. (10) and mean square error (MSE) between generated and real digits of all minority data corresponding to majority data (30000 samples) are presented in Table 2.

Based on Table 2, MSE has lower sensitivity when comparing real and generated image data, while FID better reflects the quality of generated image data concerning the imbalance ratio.

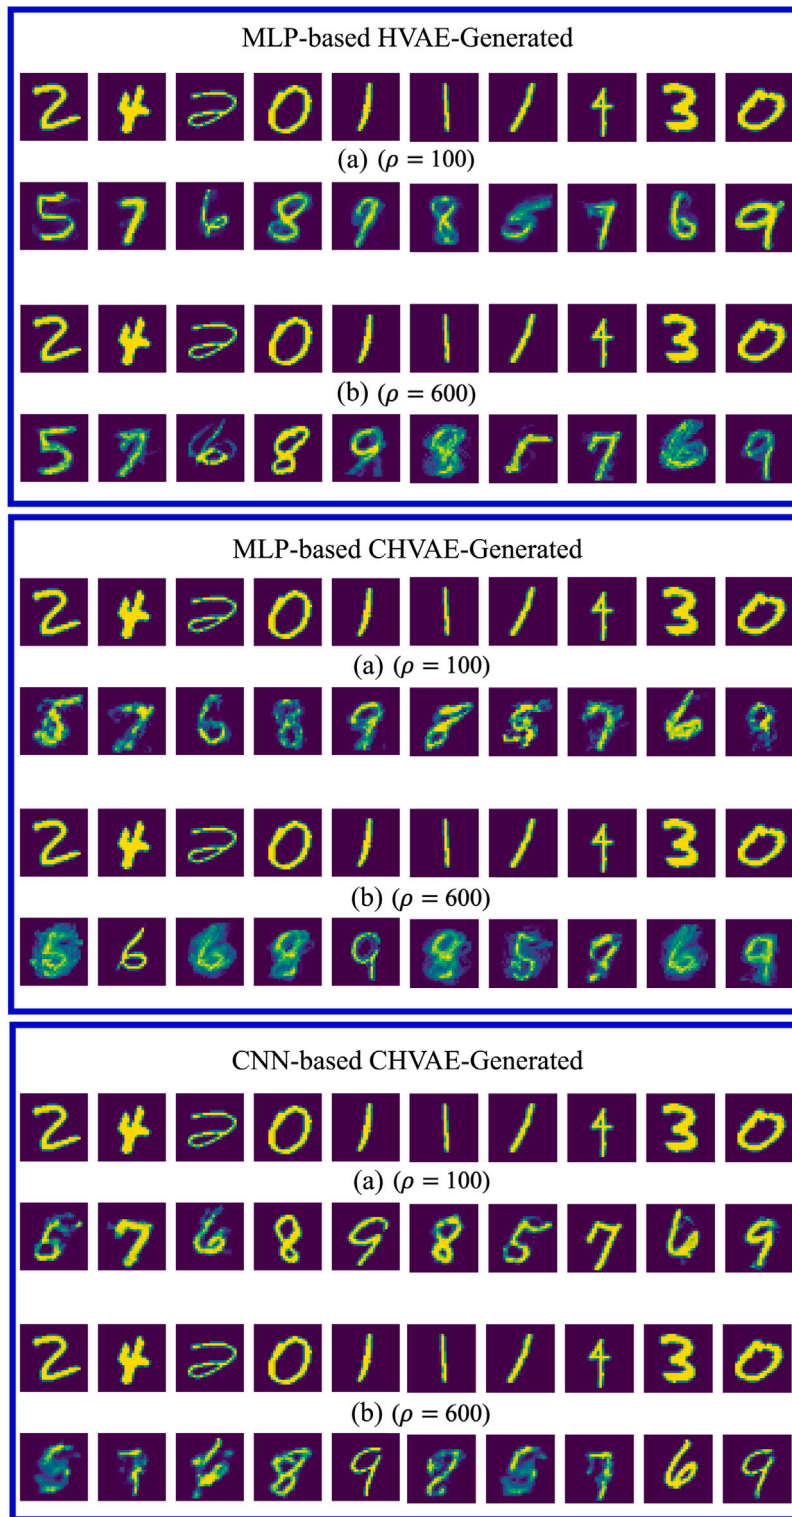


Fig. 5. Comparison of MNIST data generation results by MLP-based HVAE, MLP-based CHVAE, and CNN-based CHVAE, showing majority real data (Classes 0–4) and minority generated data (Classes 5–9) at imbalance ratios (a) $\rho = 100$, and (b) $\rho = 600$.

3.1.3. Downstream application study (behavioral assessment)

Based on Ai et al. (2023), we utilize three evaluation metrics to assess the classification performance across all methods on a balanced test set: Balanced Accuracy (B-ACC) (Huang et al., 2016; Kim et al., 2020), Average Class Specific Accuracy (ACSA) (Huang et al., 2016; Mullick et al., 2019), and Geometric Mean (GM) (Branco et al., 2017). B-ACC

corresponds to the standard accuracy metric for balanced datasets, while ACSA and GM are class-agnostic, making them more appropriate for evaluating performance under imbalanced data conditions. The numerical results of the evaluation are presented in Table 3 and illustrated in Figs. 8 and 9 representing the average of three independent trials of the classification.



Fig. 6. Visualization of MNIST data generation by CNN-based CHVAE.

In Table 3, we have compared the performance of CHVAE against the baseline and other recent oversampling approaches, which are briefly presented below:

1. **Empirical risk minimization (ERM)**: it is the conventional training approach utilizing cross-entropy loss, without any balancing procedures.
2. **Random over-sampling (ROS)** (Japkowicz, 2000): it balances the sampling distribution through repeated random sampling of the minority class.
3. **SMOTE** (Chawla et al., 2002): it balances the sampling distribution through linear interpolation of nearest neighbors within the minority class.
4. **Re-weighting (RW)** (Huang et al., 2016): it adjusts the objective function based on the size of class samples.
5. **Class-balanced re-weighting (CBRW)** (Cui et al., 2019): it is a modified version of re-weighting (RW), which introduces the effective sample number $E_k = \frac{1-\beta^{N_k}}{1-\beta}$ for each class, with β set to 0.9999.
6. **FOCAL** (Lin et al., 2017): it seeks to balance the sample-wise classification loss during model training by reducing the weight of well-classified samples.
7. **LDAM** (Cao et al., 2019): it defines a margin loss that is aware of label distribution, promoting larger margins for few-shot classes.
8. **OCVAE** (Fajardo et al., 2021): it includes the oversampling of the minority class using a Conditional VAE, which is a generative model that augments the dataset conditioned on class labels.
9. **OCGAN or OCDCGAN** (Fajardo et al., 2021): it oversamples the minority class using a Conditional GAN, which is a generative model that augments the dataset based on class labels.

10. **MGVAE** (Ai et al., 2023): it presents an objective function with Elastic Weight Consolidation (EWC) regularization that facilitates the augmentation of minority data through the utilization of a majority-based prior.

Based on Table 3, our HVAE framework significantly outperforms the other methods evaluated in Ai et al. (2023), consistently achieving superior results across all comparisons. Note that they also left out OCGAN’s results because it cannot train correctly on small-scale data. In cases of extreme imbalance ($\rho = 600$), the percentage improvement of the GM term relative to MGVAE is 10.2 for the MLP-based HVAE and 3.7 for the CNN-based CHVAE. These results show that MLP-based HVAE significantly outperforms other methods, while CNN-based CHVAE retains its effectiveness even in conditions of extreme imbalance. Additionally, the CNN-based CHVAE demonstrates significant results when $\rho = 100$. The fine-tuning monitoring indicators, computed based on Eqs. (7) and (8) and plotted during training, are illustrated in Fig. 10 (for $\rho = 100$). The reconstruction and FID losses converge to zero, as shown in Figs. 10(a) and 10(b). In other words, the fine-tuned CHVAE learns to capture the minority data structure by updating information in each epoch, thereby minimizing losses. Fig. 10(c) shows the KL loss converging to a small, non-negative value near zero, indicating effective construction of z_2 latent distribution (Eq. (6)). Building z_2 latent distribution during fine-tuning acts as a regularizer to prevent overfitting in the conditional decoder of the fine-tuned CHVAE, enabling it to generate diverse samples of minority data, even with limited examples used during fine-tuning. This diversity is demonstrated in Fig. 6 and in the results provided in Table 3.

3.2. FOWT domain translation for different mooring system

In this section, we aim to develop a real-time monitoring framework for different mooring systems in FOWTs by utilizing the CHVAE feature for domain translation. We evaluate our proposed framework using tension records simulated by a numerical model of the benchmark DeepCwind OC4 platform (Robertson et al., 2014a), depicted in Fig. 11. To assess the generalization capability of our approach, we consider two distinct mooring systems: the source system (OC4), used for training the CHVAE, and the target system (modified OC4 mooring system, with equivalent pretension), which differs structurally and is used for evaluating domain translation and classification on previously unseen configurations. This setting reflects a realistic operational context where labeled data from a new or modified mooring system may not be available, especially during the design stage of a new concept of a mooring system. The details of these systems are presented in Table 4. Additionally, we describe the necessary preliminary steps for data preparation and CHVAE implementation.

3.2.1. Data preprocessing

The tension forces in the fairleads of the semisubmersible floating system (Fig. 11) are treated as key variables for monitoring the health of the mooring system. The tension in the mooring system represents a good measure of the excitation that the mooring system is undergoing. However, direct measurements of such a variable can be possible only via load cells, whose reliability is generally low. For this reason, the scientific and technical communities are working on assessing the tension of the mooring lines using indirect measurements.

The time-series responses as the training, validating, and testing data are obtained under different randomly sampled combinations of the environmental and operational conditions, including wave-significant height (H_s) ranging from 1 to 7 m, peak period (T_p) between 8 to 15 s, wind speed (V) from 1 to 10 m per second, and current (C) from 0.5 to 1.5 m per second, which altogether account for the variability of the excitation. Please note that with no lack of generality, the wave conditions refer to a generic site. For example, in BiMEP (see Ezepeleta et al., 2025), the considered metoceanic conditions cover more

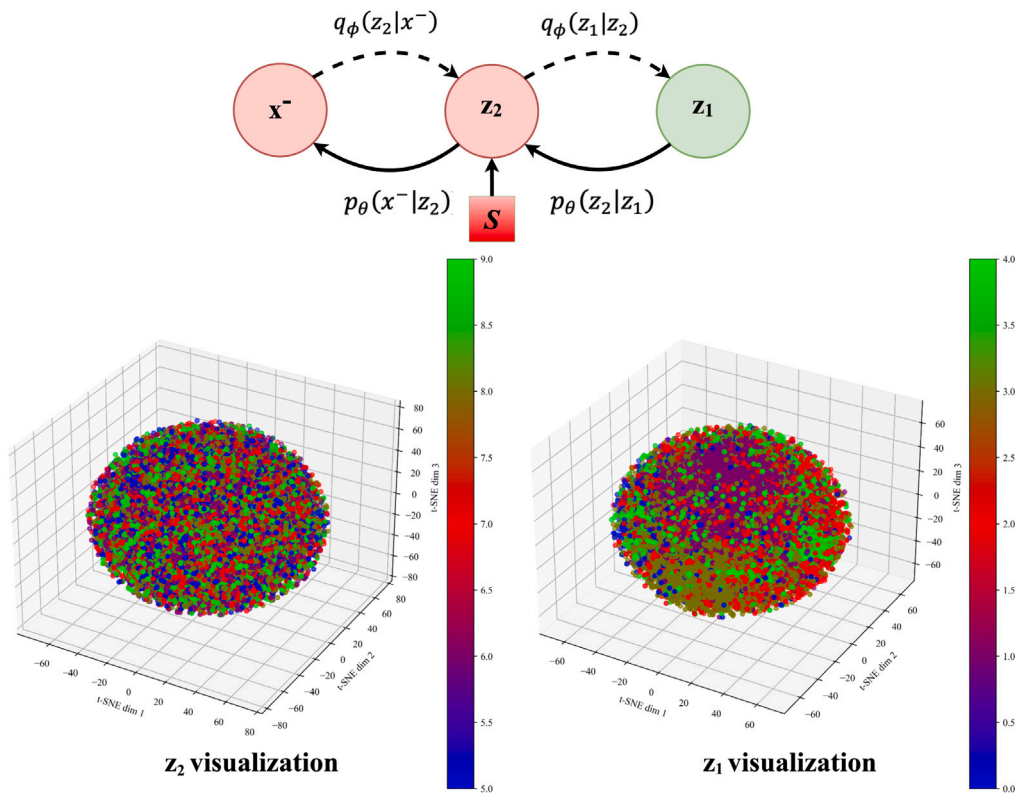


Fig. 7. The architecture of conditional hierarchical VAE concept.

Table 3

Comparison of CHVAE classification performance on MNIST under two different imbalance ratios (IR).

IR	$\rho = 100$			$\rho = 600$		
	B-ACC	ASCA	GM	B-ACC	ASCA	GM
ERM	51.4 ± 0.0	50.0 ± 0.0	0.0 ± 0.0	51.4 ± 0.0	50.0 ± 0.0	0.0 ± 0.0
FOCAL	51.4 ± 0.0	50.0 ± 0.0	0.0 ± 0.0	51.4 ± 0.0	50.0 ± 0.0	0.0 ± 0.0
RW	77.4 ± 1.2	76.7 ± 1.2	73.1 ± 1.3	59.9 ± 1.6	58.8 ± 1.4	41.7 ± 1.8
CBRW	75.1 ± 0.8	74.3 ± 0.7	69.8 ± 1.3	56.1 ± 0.5	55.1 ± 0.3	31.2 ± 1.2
LDAM	82.9 ± 0.5	82.4 ± 0.6	80.3 ± 0.7	63.1 ± 0.9	62.0 ± 0.8	48.7 ± 1.0
RS	79.2 ± 0.3	78.5 ± 0.3	75.4 ± 0.2	58.5 ± 1.0	57.3 ± 1.1	37.7 ± 2.0
SMOTE	80.6 ± 0.3	80.1 ± 0.2	77.4 ± 0.1	60.0 ± 0.9	58.3 ± 1.0	40.2 ± 1.1
OCVAE	83.0 ± 0.4	82.6 ± 0.4	80.6 ± 0.6	63.8 ± 0.2	62.8 ± 0.5	50.7 ± 0.5
MGVAE	85.0 ± 0.2	84.6 ± 0.2	83.2 ± 0.2	65.4 ± 1.0	64.4 ± 1.1	53.4 ± 1.1
HVAE w MLP	87.0 ± 0.2	86.5 ± 0.2	85.6 ± 0.2	71.9 ± 0.3	71.0 ± 0.4	63.6 ± 1.0
CHVAE w MLP	86.9 ± 0.1	86.5 ± 0.1	85.4 ± 0.2	66.1 ± 0.3	65.9 ± 0.3	55.3 ± 0.1
CHVAE w CNN	90.5 ± 0.2	90.3 ± 0.2	90.1 ± 0.2	67.3 ± 0.4	67.2 ± 0.4	57.1 ± 1.1




Fig. 8. Confusion matrix; multiclass MNIST classification (CHVAE w MLP); (a) IR = 100, (b) IR = 600.

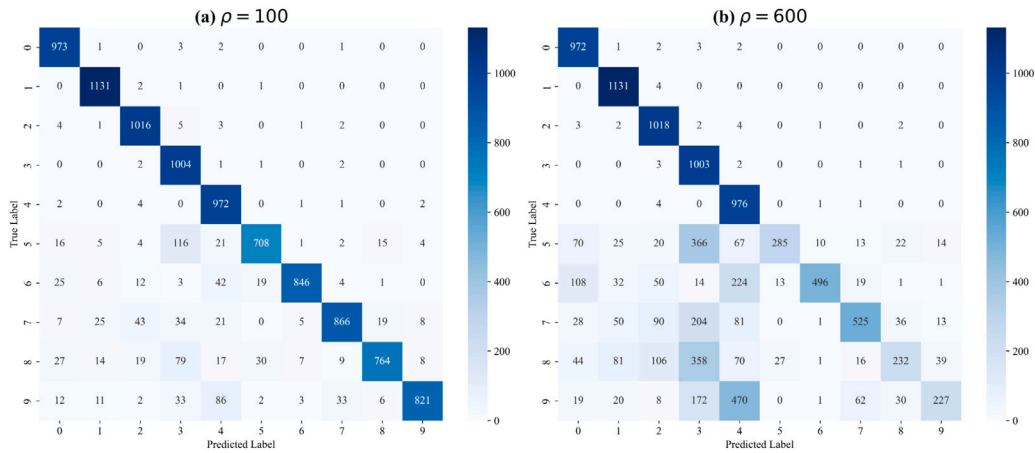


Fig. 9. Confusion matrix; multiclass MNIST classification (CHVAE w CNN); (a) IR = 100, (b) IR = 600.

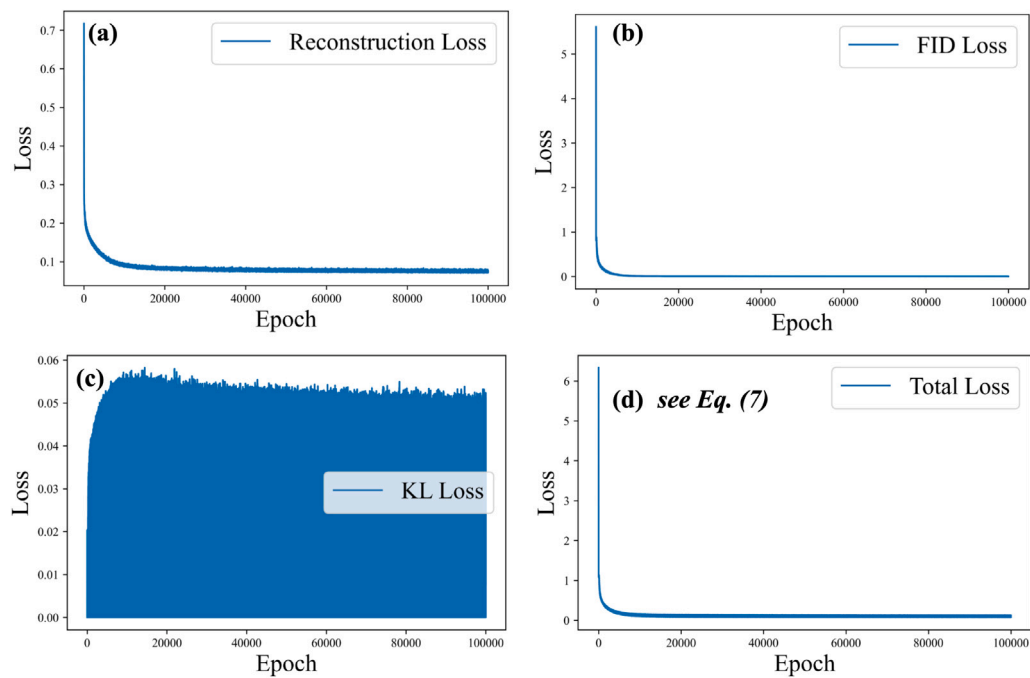


Fig. 10. Fine-tuning monitoring plots for MNIST data (CHVAE with MLP-based architecture, IR = 100).

Table 4

Properties of the mooring systems (Source and Target).

Property	Source (OC4)	Target (modified OC4)	Unit
Number of lines	3	3	-
Segments per line	20	20	-
Line diameter	7.66×10^{-2}	8.50×10^{-2}	m
Mass/length	1.13×10^2	1.44×10^2	kg/m
Axial stiffness (EA)	7.54×10^8	6.17×10^8	N
Unstretched length	8.35×10^2	8.25×10^2	m

than 80% of the yearly occurrence for the site. The simulation produces 1800 s of response data at a fixed sampling frequency of 40 Hz. The response is downsampled to 5 Hz to simulate realistic conditions, with the first 300 s skipped to exclude transient effects.

Corrosion and corrosion-related issues can be considered the primary damage mechanism affecting the mooring systems in floating structures (see, for example, Ma et al., 2013). For this reason, in this research, we have modeled the effect of the corrosion by reducing the axial stiffness (EA (N)) and mass per unit length ($Mass\ density$ (kg/m))

of all three mooring lines, each with varying degrees of degradation within the MoorDyn module. This experiment involves the degradation of mooring lines 1 and 2 by 3%, 5%, and 7%, while mooring line 3 experiences degradation of 5%, 7%, and 10% in both axial stiffness (EA) and mass per unit length, as detailed in Table 5. We focused on this type of corrosion to mimic one of the degradation scenarios that mooring lines may experience over their lifespan; based on the Orcaflex manual for the studless chains, the safety factor decreases in this case.

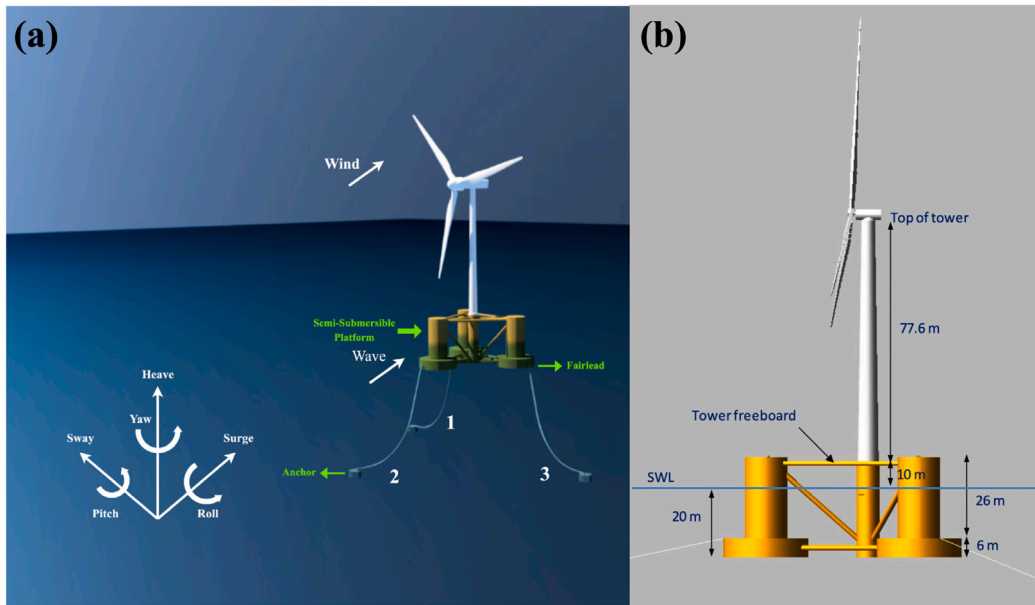


Fig. 11. OC4-DeepCwind semisubmersible FOWT; (a)Schematic sketch, (b) geometric dimensions (Robertson et al., 2014b).

Table 5

Corrosion severity levels and line-wise degradation configuration used in CHVAE training for the source mooring system.

Severity level	Lines 1 and 2	Line 3
3%–5%	↓3% Mass/length and EA	↓5% Mass/length and EA
5%–7%	↓5% Mass/length and EA	↓7% Mass/length and EA
7%–10%	↓7% Mass/length and EA	↓10% Mass/length and EA

Based on the provided details, each data sample consists of a 1500-second time series of fairlead tension, with 7500 time steps at a frequency of 5 Hz and three features: FairTen1, FairTen2, and FairTen3 (three mooring lines in Fig. 11). Through trial and error, the frame size selected for the input to the VAE layer is 50. Consequently, the input data shape is (batch size, frame size, num features, num channels) = (batch size, 50, 3, 1). The input data was normalized using a multi-step process to prepare it for training in a CNN-based model. The steps involved are as follows:

1. For each channel, the mean value was computed and subtracted to center the data around zero.
2. After subtracting the mean, the maximum absolute value for each channel was determined.
3. Each channel was then normalized by dividing the data by its maximum absolute value, ensuring that it was scaled between -1 and 1.

Subsequently, the denormalization parameters, specifically the mean and the maximum absolute value for each frame size and feature, are incorporated into d_1 and d_2 for the healthy and damaged data, respectively, which have the shape (num frames, frame size, 2 × num features).

In this experiment, the number of healthy state data samples is $N^+ = 625$. To train the pre-trained VAE on these 625 samples, 80% of the healthy data is used for training and validation (80% for training, 20% for validation), while the remaining 20% is allocated for testing.

For the source mooring system (OC4), a fine-tuning process is carried out using healthy and 40% damaged data to improve the performance of the CHVAE model. Following the fine-tuning phase, the model is then employed to synthesize damaged responses for the target mooring system, using the healthy data from the target system. In Fig. 12, we provide a visual overview of the proposed domain

translation framework. For example, Fig. 12 illustrates how tension records corresponding to 3%–5% corrosion severity can be generated for the target mooring system.

3.2.2. CHVAE specifications

The architecture of VAE’s encoder and decoder, used in the pre-training, is as follows:

Encoder: Input01 (Batch size, 50, 3, 1), Conv2D(32), 3 × Conv2D(64), MLP(32), Input02(d_1);

Decoder01: Input(z_1), 2 × Conv2DTranspose(64), Conv2DTranspose(32);

Decoder02(Denorm): Input01(z_1), Input02(d_1), MLP(64), MLP(32), MLP(d_1 size, LeakyRelu).

The architecture of VAE’s encoder and decoder, used in the fine-tuning is as follows:

Encoder: Input01 (Batch size, 50, 3, 1), Conv2D(32), 3 × Conv2D(64), MLP(32), Input02(d_1), Input03(S);

Decoder01: Input01(z_2), Input02(S), Conv2DTranspose(128), 3 × Conv2DTranspose(64), Conv2DTranspose(32);

Decoder02(Denorm): Input01(z_2), Input02(d_1), Input03(S), MLP(64), MLP(32), MLP(d_2 size, LeakyRelu).

The VAE characteristics are also mentioned in Table 6.

All deep learning models was trained on an Apple M1 Pro chip featuring a 10-core CPU and a 16-core integrated GPU, complemented by 16 GB of unified memory. The models were implemented using TensorFlow 2.11.0, utilizing the Keras API (version 2.11.0) for model development. The OpenFAST simulations were also executed on the DIPC cluster (Centre, 2024).

To enhance reproducibility and elucidate the computational cost of our proposed approach, we present the training time of the CNN-based CHVAE model for both case studies in Table 7. The findings encompass both the pretraining and fine-tuning phases. All experiments

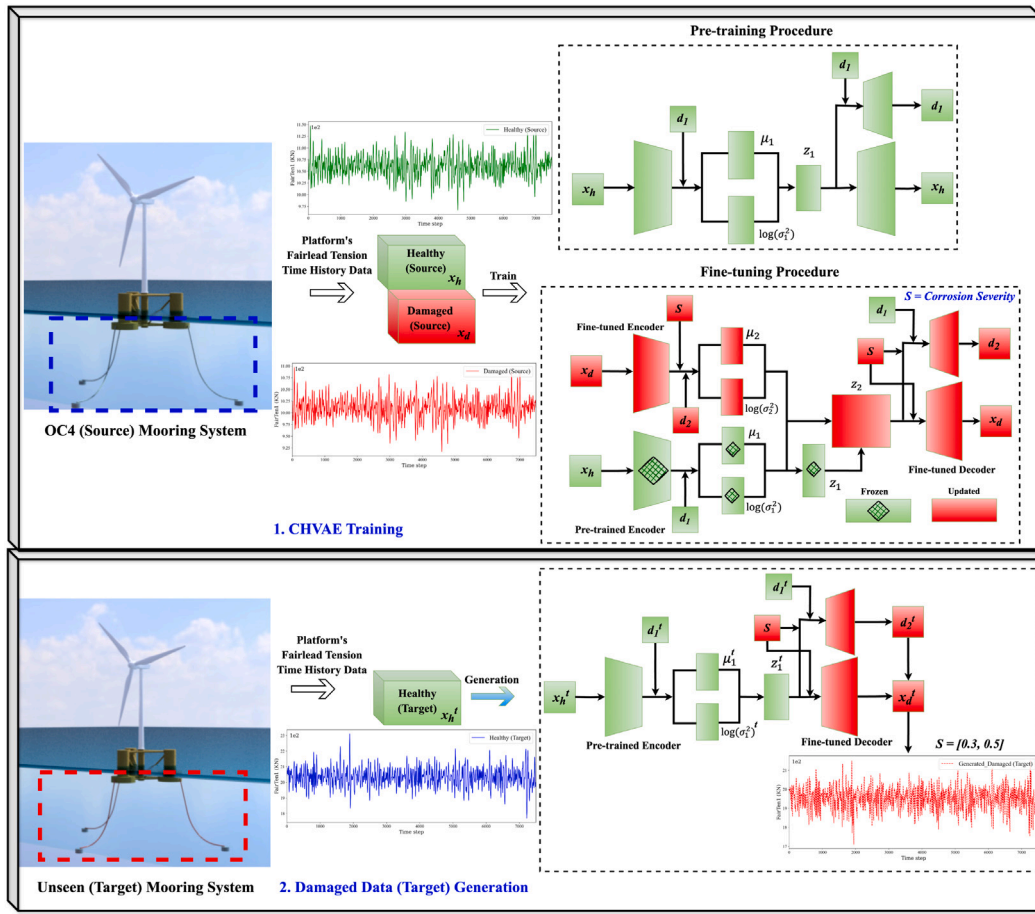


Fig. 12. Illustration of the proposed damaged data generation process for an unseen (Target) system (Domain Translation).

Table 6
The VAE parameters (pre-training and fine-tuning procedures) for FOWT data.

Parameter	Value
Batch size	pre-train:32, fine-tune: 4650
Convolution window/kernel size	2 × 2
Convolution stride	1 × 1
Padding	Same
Activation function	ReLU
Activation function (Output layer of first decoder)	Tanh
Optimizer	Adam
Learning rate	0.001
Reconstruction loss function	MSE (Eq. (8))
Latent space dimension	32

Table 7
Training times of CNN-based CHVAE model on Apple M1 Pro integrated GPU.

Case study	Epochs (Pre/Fine)	Pretraining (min)	Finetuning (min)
MNIST Data (Imbalance $\rho = 100$)	100/100,000	57.30	569.34 (9.5 h)
FOWT Data	100/1,600	38.97	249.40 (4.2 h)

were conducted locally using the integrated GPU of the M1 Pro, without the use of external accelerators.

Notably, despite a greater number of epochs during the fine-tuning stage (especially for MNIST), the training duration per epoch is typically reduced. Fine-tuning occurs on smaller, imbalanced subsets of data, with minority-class data size in each epoch, whereas pretraining utilizes the majority dataset. The computational load per epoch during fine-tuning is markedly decreased.

3.2.3. Qualitative and quantitative assessment

This section assesses the CHVAE’s performance by qualitatively comparing the generated data with the simulated data from the DeepCwind OC4 platform’s numerical model (Robertson et al., 2014b) using OpenFast. FID and reconstruction error (MSE) are explored to evaluate the resemblance between the CHVAE-generated data and the simulated data, referred to in this study as ‘real’ or ‘original’ data.

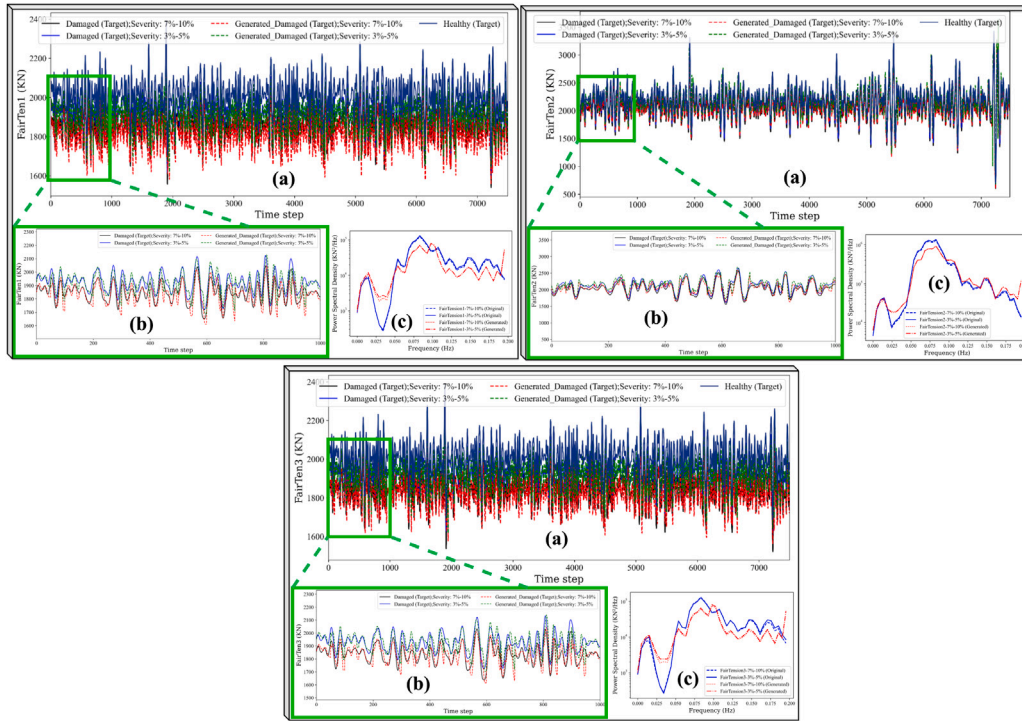


Fig. 13. Typical sample; (a) the unseen fairlead tensions (Mooring lines 1, 2, and 3) of platform coupled with target mooring system under healthy and damaged operational state when $V = 10$, $H_s = 7.0$, $T_p = 13.25$, $C = 0.5$. (b) The first 1000 time steps of real and approximated damaged record. (c) Power spectral density of tension record.

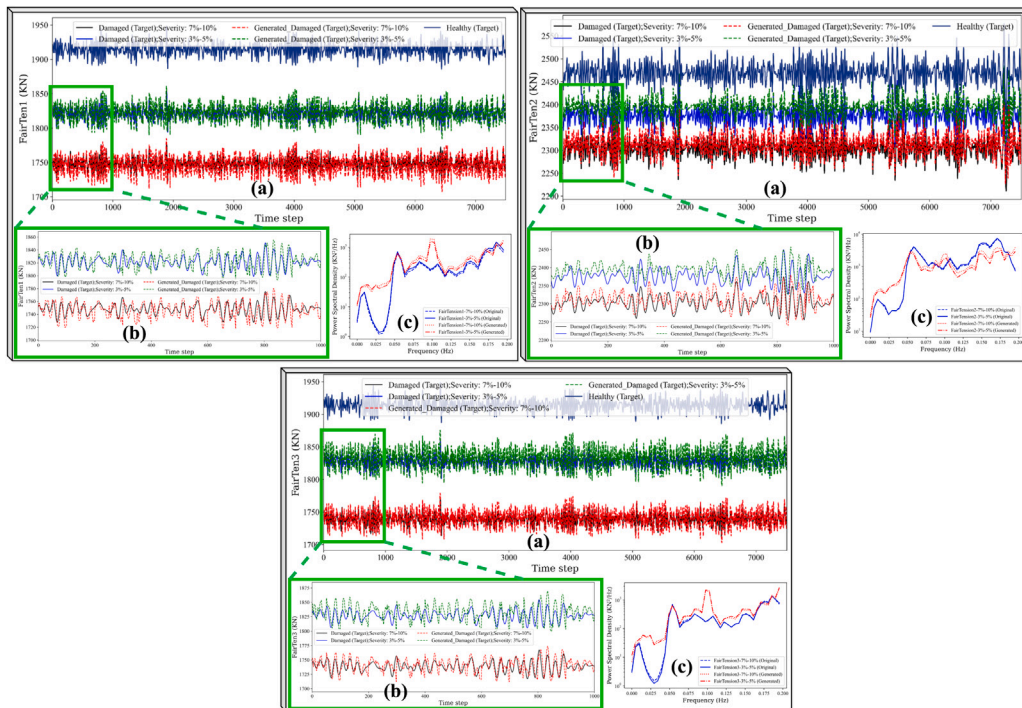


Fig. 14. Typical sample; (a) the unseen fairlead tensions (Mooring lines 1, 2, and 3) of platform coupled with target mooring system under healthy and damaged operational state when $V = 7.75$, $H_s = 1.0$, $T_p = 15.0$, $C = 1.25$. (b) The first 1000 time steps of real and approximated damaged record. (c) Power spectral density of tension record.

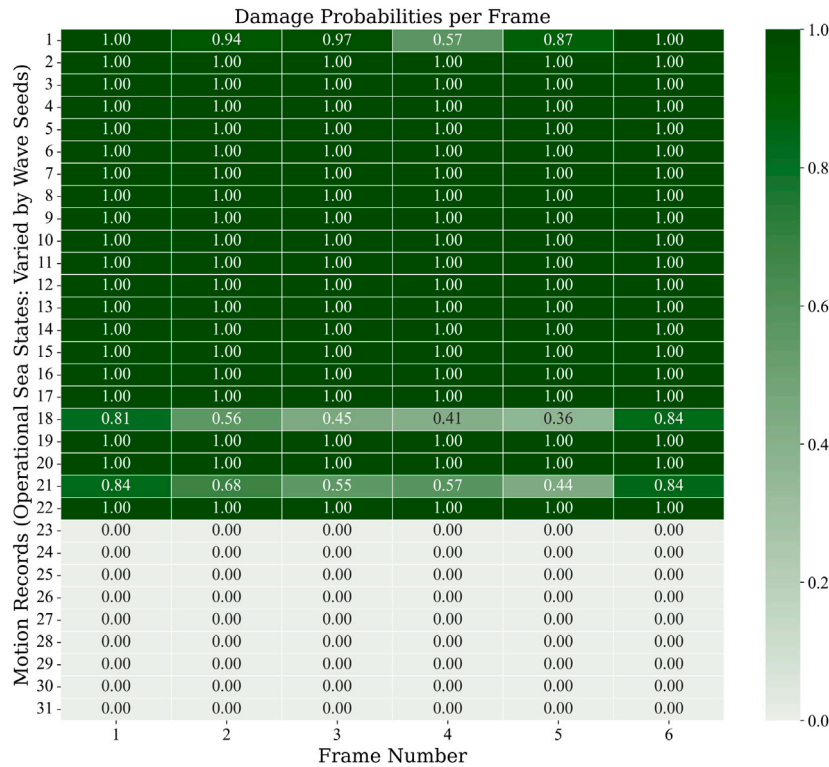


Fig. 15. Classification heat map; Trained on healthy (target) and generated damaged (target) records.

Table 8

The quality evaluation of generating damaged tension records of target mooring system by CHVAE.

Corrosion severity	FID	MSE
3%–5%	0.0725	0.4253
7%–10%	0.1314	0.4647

Visual comparisons between generated and original damaged tension records for the target mooring system, under two different, random, unseen operational sea states, are presented in Figs. 13 and 14. Two levels of corrosion severity – 3%–5% (mild) and 7%–10% (intense) – are examined in the target mooring system.

Figs. 13(a) and 14(a) show the time series of tension records for different corrosion severities as well as healthy (non-damaged) and generated damaged records. As we can see in these two figures, the discrepancies between the healthy and responses with different corrosion severities differ across various operational sea states. Consequently, CHVAE effectively captures and approximates these variations over a range of random, unseen sea states.

Figs. 13(b) and 14(b) show a zoomed-in view of the time series for a smaller segment of data. In these zoomed-in sections, the generated damaged records closely resemble the original damaged records at the specified severities, indicating that the CHVAE can generate real-scale tension patterns for different levels of damage.

Figs. 13(c) and 14(c) show the power spectral densities (PSDs) at lower frequencies, where the dominant tension dynamics occur. Comparing the generated and original records in the frequency domain, the PSDs align well across different severities, suggesting that the CHVAE preserves the spectral characteristics of the original data.

The statistical index, FID, and MSE (reconstruction error) between generated and original tension records under 625 different operational sea states for three fairleads moored with the target system are presented in Table 8.

Table 9

Corrosion severity intervals used in downstream classification, including CHVAE-generated training data and real unseen test data to evaluate generalization.

Corrosion severity interval	Description
3%–5%	CHVAE-generated training data
5%–7%	CHVAE-generated training data
7%–10%	CHVAE-generated training data
1–10%–1%–10%	Real unseen data for generalization

3.2.4. Downstream application study (behavioral assessment)

To assess the effectiveness of the generated damaged responses, a binary classification experiment was conducted. The classifier model was trained using a dataset combining healthy and generated damaged data from the target moored platform. In the testing phase, the model was evaluated using original data from unseen damage severity and random sea states, enabling real-time classification of healthy and damaged responses.

The classifier architecture is a fully connected neural network, structured as follows: $C_1 = \text{MLP} [\text{Input dim} - 256]$, $C_2 = \text{MLP} [256 - 128]$, and $C_3 = \text{MLP} [128 - 1]$. The activation function for the final dense layer is set to sigmoid, with binary cross-entropy as the loss function, optimized using the Adam optimizer with a learning rate of 1×10^{-3} . The sigmoid output represents a prediction score between 0 (healthy) and 1 (damaged).

In the training process, 80% of the responses are used, which are divided into 80% for training and 20% for validation, with the remaining 20% of the data allocated for testing. The corrosion severity intervals utilized for preparing training and testing data for subsequent binary classification are detailed in Table 9.

The input shape for each record is (number of frames, frame size \times num features) = (6, 1250×3), which yields the best performance. The classifier is trained for 25 epochs with a batch size of 200. For feature normalization, Robust Scaling is applied to each of the three features,

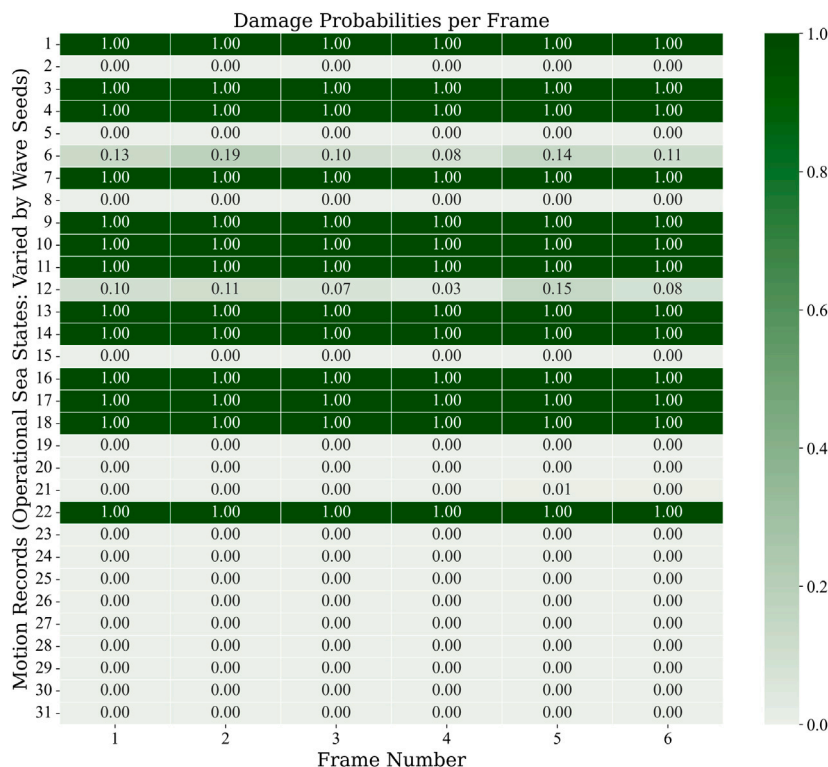


Fig. 16. Classification heat map; Trained on healthy (target) and damaged (source) records.

Table 10

The classification metrics.

Classification type	Accuracy	Precision	Recall	AUC (PR)
With Blind Domain Translation	0.9408	1.00	0.9167	0.9796
Without Blind Domain Translation	0.7096	1.00	0.5909	0.8446

as it produces the best results, particularly in handling outliers and thereby improving classification performance.

To evaluate the classifier’s performance, a set of 31 unseen tension records was selected, encompassing a range of unseen damage severities and operational sea states with different wave seeds; 22 records correspond to damaged states, while 9 pertain to healthy conditions. The classification outcomes are summarized in Figs. 15 and 16, and Table 10.

As shown in Fig. 15, unseen random damaged records are generally well identified. Even records under states numbers 18 and 21 can be considered damaged, as at least two of their frames express high probabilities of damage. Conversely, Fig. 16 demonstrates the low classifier’s performance when it lacks information about the adapted features of the target damaged records, resulting in its inability to identify real damaged records correctly.

Table 10 illustrates that the recall of the classifier increases from 0.5909 to 0.9167 with the application of blind domain translation, indicating a 32% improvement. This result demonstrates that CHVAE-based domain translation markedly improves the classifier’s capacity to identify damaged frames, thereby reinforcing the conclusion about performance enhancement across various mooring systems. This result facilitates a zero-shot classification framework, utilizing solely healthy data from the target mooring system, while the necessary damaged data is synthetically produced through the proposed CHVAE-based DGM.

4. Conclusions

This paper presents the conditional hierarchical variational autoencoder (CHVAE), a novel data generation model grounded in the

pretrain-finetune training framework. CHVAE aims to augment damage-associated data across random operational sea states in the source mooring system and synthesize labeled damaged records for the target mooring system through domain translation. This method strengthens the robustness of damage detection systems under diverse operational conditions and facilitates knowledge transfer between different mooring systems. The study also demonstrates that the behavior of the generated records in downstream applications allows a classifier with a shallow architecture to accurately label unseen, real-time data from the target mooring system.

The training and evaluation data were derived from simulations of the DeepCWind floater conducted using the OpenFast open-source platform. Moreover, the performance of the proposed data generation model is evaluated on the MNIST benchmark image dataset and compared with conventional and recently proposed oversampling methods. The following conclusions can be summarized:

- Introducing conditional layers to the encoder and decoder of fine-tuned CHVAE, besides the wisely controlling oversampling, enables the generation of diverse minority classes (e.g., digits or damage severity) from a specific majority class (Healthy) (one-to-all capability).
- The CNN-based CHVAE in MNIST multi-class classification significantly outperforms other oversampling methods at $\rho = 100$ and in scenarios of extreme imbalance at $\rho = 600$. The geometric mean index for CNN-based CHVAE increases by 3.7%, while the MLP-based HVAE shows an improvement of 10.2%.
- By incorporating the KL loss term into the CHVAE fine-tuning loss function and proposing a novel reparameterization approach

to reverse the diffusion process, we prevent overfitting, even in extreme scarcity of minority data.

- Equipping CHVAE with a second decoder to estimate de-normalization parameters enables our model to mimic real-scale time-series tension records with acceptable accuracy.
- For different mooring systems in FOWT, examining the domain translation feature of CHVAE improves the performance of a real-time, shallow-architecture zero-shot damage identification model by up to 32%.

CRedit authorship contribution statement

Hamed Fathnejat: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Vincenzo Nava:** Writing – review & editing, Validation, Supervision, Resources, Project administration, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research is financially supported by the Spanish Ministry of Economic Affairs and Digital Transformation through the Recovery, Transformation, and Resilience Plan, specifically in the R&D Missions within the Artificial Intelligence 2021 Programme. The funding is allocated within the framework of the IA4TES project (Artificial Intelligence for Sustainable Energy Transition) under reference number MIA.2021.M04.008. The authors also would like to acknowledge the ELKARTEK project RUL-ET funded by the Basque Government (KK-2024/00086); the “BCAM Severo Ochoa” accreditation of excellence CEX2021-001142-S/MICIN / AEI/10.13039/501100011033; and the Basque Government through the BERC 2022–2025 program.

Data availability

Data will be made available on request.

References

- Ahmadi-Nedushan, B., Fathnejat, H., 2022. A modified teaching–learning optimization algorithm for structural damage detection using a novel damage index based on modal flexibility and strain energy under environmental variations. *Eng. Comput.* 38 (1), 847–874. <http://dx.doi.org/10.1007/s00366-020-01197-3>.
- Ai, Q., Wang, P., He, L., Wen, L., Pan, L., Xu, Z., 2023. Generative Oversampling for Imbalanced Data via Majority-Guided VAE. *Proc. Mach. Learn. Res.* 206, 3315–3330, URL <https://arxiv.org/abs/2302.10910v1>.
- Altabay, W.A., Wu, Z., Noori, M., Fathnejat, H., 2023. Structural Health Monitoring of Composite Pipelines Utilizing Fiber Optic Sensors and an AI-Based Algorithm—A Comprehensive Numerical Study. *Sensors* 23 (8), 3887. <http://dx.doi.org/10.3390/s23083887>, URL <https://www.mdpi.com/1424-8220/23/8/3887>.
- Branco, P., Torgo, L., Ribeiro, R.P., 2017. A Survey of Predictive Modeling on Imbalanced Domains. *ACM Comput. Surv.* 49 (2), 1–50. <http://dx.doi.org/10.1145/2907070>, URL <https://dl.acm.org/doi/10.1145/2907070>.
- Cao, K., Wei, C., Gaidon, A., Arechiga, N., Ma, T., 2019. Learning imbalanced datasets with label-distribution-aware margin loss. *Adv. Neural Inf. Process. Syst.* 32.
- Centre, D.I.P., 2024. Atlas-EDR, Technical Documentation. URL <https://dipc.ehu.es/en>.
- Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P., 2002. SMOTE: Synthetic Minority Over-sampling Technique. *J. Artificial Intelligence Res.* 16, 321–357. <http://dx.doi.org/10.1613/JAIR.953>, URL <https://www.jair.org/index.php/jair/article/view/10302>.
- Chen, J., Kim, M.H., 2021. Review of Recent Offshore Wind Turbine Research and Optimization Methodologies in Their Design. *J. Mar. Sci. Eng.* 10 (1), 28. <http://dx.doi.org/10.3390/jmse10010028>, URL <https://www.mdpi.com/2077-1312/10/1/28>.

- Chen, P., Zhao, R., He, T., Wei, K., Yuan, J., 2023. Unsupervised structure subdomain adaptation based the Contrastive Cluster Center for bearing fault diagnosis. *Eng. Appl. Artif. Intell.* 122, 106141. <http://dx.doi.org/10.1016/J.ENGAPPAL.2023.106141>.
- Coraga, E.M., Ferreira, J.V., Nóbrega, E.G., 2023. An unsupervised structural health monitoring framework based on Variational Autoencoders and Hidden Markov Models. *Reliab. Eng. Syst. Saf.* 231, 109025. <http://dx.doi.org/10.1016/j.res.2022.109025>, URL <https://linkinghub.elsevier.com/retrieve/pii/S0951832022006408>.
- Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S., 2019. Class-balanced loss based on effective number of samples. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9268–9277.
- DNV, 2022. Floating Offshore Wind: The Next Five Years. URL <https://www.dnv.com/focus-areas/floating-offshore-wind/floating-offshore-wind-the-next-five-years/>.
- Edwards, E.C., Holcombe, A., Brown, S., Ransley, E., Hann, M., Greaves, D., 2023. Evolution of floating offshore wind platforms: A review of at-sea devices. *Renew. Sustain. Energy Rev.* 183, 113416. <http://dx.doi.org/10.1016/j.rser.2023.113416>, URL <https://linkinghub.elsevier.com/retrieve/pii/S1364032123002733>.
- Ezpeleta, H., Ulazia, A., Aristondo, O., Ibarra-Berastegi, G., 2025. Future wave energy evolution in the Bay of Biscay for different converters. *Energy* 331, 136713. <http://dx.doi.org/10.1016/j.energy.2025.136713>.
- Fajardo, V.A., Findlay, D., Jaiswal, C., Yin, X., Houmanfar, R., Xie, H., Liang, J., She, X., Emerson, D., 2021. On oversampling imbalanced data with deep conditional generative models. *Expert Syst. Appl.* 169, 114463. <http://dx.doi.org/10.1016/j.eswa.2020.114463>, URL <https://linkinghub.elsevier.com/retrieve/pii/S0957417420311155>.
- Fathnejat, H., Ahmadi-Nedushan, B., 2017. Structural Damage Detection by Sensitivity-based method and Cascade Feed-forward Neural Network Based on Proper Orthogonal Modes. In: *6th National and 2nd International Conference on New Materials and Structures in Civil Engineering*, Yazd, Iran, 2017.. URL https://www.researchgate.net/publication/325379505_Structural_Damage_Detection_by_Sensitivity-based_method_and_Cascade_Feed-forward_Neural_Network_Based_on_Proper_Orthogonal_Modes.
- Fathnejat, H., Ahmadi-Nedushan, B., 2020. An efficient two-stage approach for structural damage detection using meta-heuristic algorithms and group method of data handling surrogate model. *Front. Struct. Civ. Eng.* 14 (4), 907–929. <http://dx.doi.org/10.1007/s11709-020-0628-1>, URL <https://link.springer.com/10.1007/s11709-020-0628-1>.
- Fathnejat, H., Ahmadi-Nedushan, B., Hosseinijad, S., Noori, M., Altabay, W.A., 2023. A data-driven structural damage identification approach using deep convolutional-attention-recurrent neural architecture under environmental effects. *Eng. Struct.* 276, 1–23. <http://dx.doi.org/10.1016/j.engstruct.2022.115311>, URL <https://linkinghub.elsevier.com/retrieve/pii/S0141029622013876>.
- Fathnejat, H., Nava, V., 2025. Data augmentation for damaged scenarios in floating offshore wind turbines: an approach based on diffusion architecture, hierarchical variational approximation and healthy data distribution. *Eng. Comput.* (0123456789), <http://dx.doi.org/10.1007/s00366-025-02148-6>, DOI: 10.1007/s00366-025-02148-6 <https://link.springer.com/10.1007/s00366-025-02148-6>.
- Fathnejat, H., Torkzadeh, P., Salajegheh, E., Ghiasi, R., 2014. Structural damage detection by model updating method based on cascade feed-forward neural network as an efficient approximation mechanism. *Int. J. Optim. Civ. Eng.* 4 (4), 451–472.
- Fu, H., Li, C., Liu, X., Gao, J., Celikyilmaz, A., Carin, L., 2019. Cyclical Annealing Schedule: A Simple Approach to Mitigating. In: *Proceedings of the 2019 Conference of the North. Association for Computational Linguistics*, Stroudsburg, PA, USA, pp. 240–250. <http://dx.doi.org/10.18653/v1/N19-1021>, URL <http://aclweb.org/anthology/N19-1021>.
- Ghiasi, R., Fathnejat, H., Torkzadeh, P., 2019. A three-stage damage detection method for large-scale space structures using forward substructuring approach and enhanced bat optimization algorithm. *Eng. Comput.* 35 (3), 857–874. <http://dx.doi.org/10.1007/s00366-018-0636-0>, URL <http://link.springer.com/10.1007/s00366-018-0636-0>.
- Ghiasi, R., Lestoille, N., Diaine, C., Malekjafarian, A., 2025. Unsupervised domain adaptation for drive-by condition monitoring of multiple railway tracks. *Eng. Appl. Artif. Intell.* 139, 109516. <http://dx.doi.org/10.1016/J.ENGAPPAL.2024.109516>.
- Gigliotti, V., Poole, J., Venanzi, I., Ubertini, F., Worden, K., 2024. A domain adaptation approach to damage classification with an application to bridge monitoring. *Mech. Syst. Signal Process.* 209, <http://dx.doi.org/10.1016/j.ymsp.2024.111135>.
- Gorostidi, N., Pardo, D., Nava, V., 2023. Diagnosis of the health status of mooring systems for floating offshore wind turbines using autoencoders. *Ocean Eng.* 287, 115862. <http://dx.doi.org/10.1016/j.oceaneng.2023.115862>, URL <https://linkinghub.elsevier.com/retrieve/pii/S0029801823022461>.
- Havtorn, J.D., Frellsen, J., Hauberg, S., Maaløe, L., 2021. Hierarchical VAEs Know What They Don't Know. Technical Report.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S., 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. *Adv. Neural Inf. Process. Syst.* 30.
- Ho, J., Jain, A., Abbeel, P., 2020. Denoising Diffusion Probabilistic Models. In: *Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (Eds.), In: Advances in Neural Information Processing Systems*, vol. 33, Curran Associates, Inc., pp. 6840–6851, URL https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bfcfc8584af0d967f1ab10179ca4b-Paper.pdf.

- Huang, C., Li, Y., Loy, C.C., Tang, X., 2016. Learning deep representation for imbalanced classification. *Proc. IEEE Comput. Soc. Conf. Computer Vis. Pattern Recognit.* 2016-December, 5375–5384. <http://dx.doi.org/10.1109/CVPR.2016.580>.
- Japkowicz, N., 2000. The class imbalance problem: Significance and strategies. In: *Proc. of the Int'l Conf. on Artificial Intelligence*, vol. 56, pp. 111–117.
- Jiang, X., Day, S., Clelland, D., Liang, X., 2019. Analysis and real-time prediction of the full-scale thrust for floating wind turbine based on artificial intelligence. *Ocean Eng.* 175, 207–216. <http://dx.doi.org/10.1016/j.oceaneng.2019.01.046>.
- Jiang, X., Li, Z., Tao, L., Chen, Y., Zeng, W., Cai, C., Zhu, G., Li, Q., 2024. Development and characterisation of an AI-in-the-loop testing platform for floating wind turbines PART I: Construction, validation, and benchmark testing. *Ocean Eng.* 297, 116968. <http://dx.doi.org/10.1016/j.oceaneng.2024.116968>.
- Kim, J., Jeong, J., Shin, J., 2020. M2m: Imbalanced Classification via Major-to-Minor Translation. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition. CVPR, IEEE, pp. 13893–13902. <http://dx.doi.org/10.1109/CVPR42600.2020.01391>, URL <https://ieeexplore.ieee.org/document/9156943/>.
- Kiranyaz, S., Devecioglu, O.C., Alhams, A., Sassi, S., Ince, T., Abdeljaber, O., Avci, O., Gabbouj, M., 2024. Zero-shot motor health monitoring by blind domain transition. *Mech. Syst. Signal Process.* 210, 111147. <http://dx.doi.org/10.1016/j.ymssp.2024.111147>, URL <https://linkinghub.elsevier.com/retrieve/pii/S0888327024000451>.
- Lecun, Y., Bottou, L., Bengio, Y., Haffner, P., 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86 (11), 2278–2324. <http://dx.doi.org/10.1109/5.726791>, URL <http://ieeexplore.ieee.org/document/726791/>.
- Lee, S., Jang, K., Lee, S., Cho, H., Shin, S., 2023. Parametric model order reduction by machine learning for fluid–structure interaction analysis. *Eng. Comput.* 1, 3. <http://dx.doi.org/10.1007/s00366-023-01782-2>, URL <https://link.springer.com/10.1007/s00366-023-01782-2>.
- Li, L., Betti, R., 2023. A machine learning-based data augmentation strategy for structural damage classification in civil infrastructure system. *J. Civ. Struct. Heal. Monit.* 13 (6–7), 1265–1285. <http://dx.doi.org/10.1007/s13349-023-00705-5>, URL <https://link.springer.com/10.1007/s13349-023-00705-5>.
- Li, W., Liu, D., Li, Y., Hou, M., Liu, J., Zhao, Z., Guo, A., Zhao, H., Deng, W., 2024. Fault diagnosis using variational autoencoder GAN and focal loss CNN under unbalanced data. *Struct. Heal. Monit.* <http://dx.doi.org/10.1177/14759217241254121>, URL <https://journals.sagepub.com/doi/10.1177/14759217241254121>.
- Lin, Y.Z., Nie, Z.H., Ma, H.W., 2017. Structural Damage Detection with Automatic Feature-Extraction through Deep Learning. *Computer-Aided Civ. Infrastruct. Eng.* 32 (12), 1025–1046. <http://dx.doi.org/10.1111/mice.12313>.
- Longman, R.P., Xu, Y., Sun, Q., Turkan, Y., Riggio, M., 2023. Digital Twin for Monitoring In-Service Performance of Post-Tensioned Self-Centering Cross-Laminated Timber Shear Walls. *J. Comput. Civ. Eng.* 37 (2), 04022055. [http://dx.doi.org/10.1061/\(ASCE\)CP.1943-5487.0001050](http://dx.doi.org/10.1061/(ASCE)CP.1943-5487.0001050), URL <https://ascelibrary.org/doi/10.1061/%28ASCE%29CP.1943-5487.0001050>.
- Ma, X., Lin, Y., Nie, Z., Ma, H., 2020. Structural damage identification based on unsupervised feature-extraction via Variational Auto-encoder. *Measurement* 160, 107811. <http://dx.doi.org/10.1016/J.MEASUREMENT.2020.107811>.
- Ma, K.T., Shu, H., Smedley, P., L'Hostis, D., Duggal, A., 2013. A Historical Review on Integrity Issues of Permanent Mooring Systems. In: *All Days. OTC*, <http://dx.doi.org/10.4043/24025-MS>, URL <https://onpetro.org/OTCONF/proceedings/13OTC/All-13OTC/Houston,%20Texas,%20USA/37626>.
- Medina-Manuel, A., Molina Sánchez, R., Souto-Iglesias, A., 2024. AI-Driven Model Prediction of Motions and Mooring Loads of a Spar Floating Wind Turbine in Waves and Wind. *J. Mar. Sci. Eng.* 12 (9), 1464. <http://dx.doi.org/10.3390/jmse12091464>.
- Mullick, S.S., Datta, S., Das, S., 2019. Generative Adversarial Minority Oversampling. In: 2019 IEEE/CVF International Conference on Computer Vision. ICCV, 2019-October, IEEE, pp. 1695–1704. <http://dx.doi.org/10.1109/ICCV.2019.00178>, URL <https://ieeexplore.ieee.org/document/9008836/>.
- Mylonas, C., Abdallah, I., Chatzi, E.N., 2020. Deep Unsupervised Learning for Condition Monitoring and Prediction of High Dimensional Data with Application on Windfarm SCADA Data. In: *Conference Proceedings of the Society for Experimental Mechanics Series*. Springer, pp. 189–196. http://dx.doi.org/10.1007/978-3-030-12075-7_21, URL http://link.springer.com/10.1007/978-3-030-12075-7_21.
- M'zoughi, F., Garrido, I., Garrido, A.J., De La Sen, M., 2024a. Metaheuristic Airflow control for vibration mitigation of a hybrid oscillating water Column-Floating offshore wind turbine system. *Energy Convers. Manag.: X* 23, 100629. <http://dx.doi.org/10.1016/j.ecmx.2024.100629>.
- M'zoughi, F., Lekube, J., Garrido, A.J., De La Sen, M., Garrido, I., 2024b. Machine learning-based diagnosis in wave power plants for cost reduction using real measured experimental data: Mutriku Wave Power Plant. *Ocean Eng.* 293, 116619. <http://dx.doi.org/10.1016/j.oceaneng.2023.116619>.
- Nava, V., Ruiz-Minguela, P., Perez-Moran, G., Rodriguez-Arias, R., Lopez-Mendia, J., Villate-Martinez, J.L., 2019. Installation, operation and maintenance of offshore renewables. In: *Renewable Energy from the Oceans: From Wave, Tidal and Gradient Systems To Offshore Wind and Solar*. Institution of Engineering and Technology, pp. 397–424. <http://dx.doi.org/10.1049/PBPO129Ech11>.
- Omid, S., Liang, S.X., 2019. Vibration-based semantic damage segmentation for large-scale structural health monitoring. <http://dx.doi.org/10.1111/mice.12523>, URL <https://onlinelibrary.wiley.com/doi/10.1111/mice.12523>.
- OpenFAST, 2023. OpenFAST Documentation. URL <https://openfast.readthedocs.io/en/main/>.
- Orcina, 2022. Chain: OrcaFlex User Manual. Version 12.1. URL <http://www.orcina.com/webhelp/OrcaFlex/Content/html/Chain.htm>.
- Pollastro, A., Testa, G., Bilotta, A., Preverte, R., 2022. Semi-supervised detection of structural damage using Variational Autoencoder and a One-Class Support Vector Machine. URL <http://arxiv.org/abs/2210.05674>.
- Ragab, M., Chen, Z., Wu, M., Li, H., Kwoh, C.-K., Yan, R., Li, X., 2021. Adversarial Multiple-Target Domain Adaptation for Fault Classification. *IEEE Trans. Instrum. Meas.* 70, 1–11. <http://dx.doi.org/10.1109/TIM.2020.3009341>, URL <https://ieeexplore.ieee.org/document/9141312/>.
- Rezaniaiee Aqdam, H., Etefagh, M.M., Hassannejad, R., 2018. Health monitoring of mooring lines in floating structures using artificial neural networks. *Ocean Eng.* 164, 284–297. <http://dx.doi.org/10.1016/j.oceaneng.2018.06.056>, URL <https://linkinghub.elsevier.com/retrieve/pii/S0029801818311375>.
- Robertson, A., Jonkman, J., Masciola, M., Song, H., 2014a. Definition of the semisubmersible floating system for phase II of OC4. URL <https://www.osti.gov/biblio/1155123>.
- Robertson, A., Jonkman, J., Masciola, M., Song, H., Goupee, A., Coulling, A., Luan, C., 2014b. Definition of the Semisubmersible Floating System for Phase II of OC4.
- Sharma, S., Nava, V., 2024. Condition monitoring of mooring systems for Floating Offshore Wind Turbines using Convolutional Neural Network framework coupled with Autoregressive coefficients. *Ocean Eng.* 302, 117650. <http://dx.doi.org/10.1016/J.OCEANENG.2024.117650>.
- Sohl-Dickstein, J., Weiss, E.A., Maheswaranathan, N., Ganguli, S., 2015. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. URL <http://arxiv.org/abs/1503.03585>.
- Tamuly, P., Sharma, S., Nava, V., 2025. Integrated damage detection and time-series data augmentation for floating offshore mooring systems via variational semi-supervised learning. *Ocean Eng.* 330, 121199. <http://dx.doi.org/10.1016/j.oceaneng.2025.121199>, URL <https://linkinghub.elsevier.com/retrieve/pii/S0029801825009126>.
- Torkzadeh, P., Fathnejat, H., Ghiasi, R., 2016. Damage detection of plate-like structures using intelligent surrogate model. *Smart Struct. Syst.* 18 (6), 1233–1250. <http://dx.doi.org/10.12989/sss.2016.18.6.1233>.
- Van Der Maaten, L., Hinton, G., 2008. Visualizing Data using t-SNE. *J. Mach. Learn. Res.* 9, 2579–2605.
- Williams, R., Zhao, F., Backwell, B., Lee, J., Patel, A., Hutchinson, M., Qiao, L., Nguyen, T., Lathigara, A., Liang, W., Fang, E., Cheong, J., Tan, H., Francisco, A.M., Bui, V., Nguyen, T., Fiestas, R., Cox, R., Ruas, M., Pontes De Lima, J., Rabie, H., Gitobu, J., Ladwa, R., Madan, K., Shardul, M., Muchiri, W., Weekes, N., Ogilvie, C., Khinda, N., Yamamura, J., Stedman Mhi, T., Wind, V.O., 2024. GWECs Global Offshore Wind Report - Global Wind Energy Council. Technical Report, URL <https://gwec.net/strong-2023-offshore-wind-growth-as-industry-sets-course-for-record-breaking-decade/>.
- Zhao, X., Yao, J., Deng, W., Jia, M., Liu, Z., 2022. Normalized Conditional Variational Auto-Encoder with adaptive Focal loss for imbalanced fault diagnosis of Bearing-Rotor system. *Mech. Syst. Signal Process.* 170, <http://dx.doi.org/10.1016/j.ymssp.2022.108826>.