


Using graph neural networks and frequency domain data for automated operational modal analysis of populations of structures

Journal Article

Author(s):

Jian, Xudong ; Xia , Yutong; Duthé, Gregory ; Bacsa , Kiran; Liu , Wei; Chatzi, Eleni 

Publication date:

2025

Permanent link:

<https://doi.org/https://doi.org/10.3929/ethz-c-000786591>

Rights / license:

[Creative Commons Attribution 4.0 International](#)

Originally published in:

Data-Centric Engineering 6, <https://doi.org/10.1017/dce.2025.10023>

RESEARCH ARTICLE 

Using graph neural networks and frequency domain data for automated operational modal analysis of populations of structures

Xudong Jian¹ , Yutong Xia², Gregory Duthé³ , Kiran Bacsa^{1,3} , Wei Liu^{1,4} and Eleni Chatzi^{1,3} 

¹Future Resilient Systems, Singapore-ETH Centre, Singapore, Singapore

²Institute of Data Science, National University of Singapore, Singapore, Singapore

³Department of Civil, Environmental and Geomatic Engineering, ETH Zurich, Zurich, Switzerland

⁴Department of Industrial Systems Engineering and Management, National University of Singapore, Singapore, Singapore

Corresponding author: Eleni Chatzi; Email: chatzi@ibk.baug.ethz.ch

Received: 09 July 2024; **Revised:** 26 July 2025; **Accepted:** 04 August 2025

Keywords: deep learning; feature propagation; graph neural network; operational modal identification; population-based structural health monitoring

Abstract

The population-based structural health monitoring paradigm has recently emerged as a promising approach to enhance data-driven assessment of engineering structures by facilitating transfer learning between structures with some degree of similarity. In this work, we apply this concept to the automated modal identification of structural systems. We introduce a graph neural network (GNN)-based deep learning scheme to identify modal properties, including natural frequencies, damping ratios, and mode shapes of engineering structures based on the power spectral density of spatially sparse vibration measurements. Systematic numerical experiments are conducted to evaluate the proposed model, employing two distinct truss populations that possess similar topological characteristics but varying geometric (size and shape) and material (stiffness) properties. The results demonstrate that, once trained, the proposed GNN-based model can identify modal properties of unseen structures within the same structural population with good efficiency and acceptable accuracy, even in the presence of measurement noise and sparse measurement locations. The GNN-based model exhibits advantages over the classic frequency domain decomposition method in terms of identification speed, as well as against an alternate multilayer perceptron architecture in terms of identification accuracy, rendering this a promising tool for PBSHM purposes.

Impact Statement

We have developed a graph neural network (GNN)-based scheme for automated operational modal analysis. Numerical experiments demonstrate that the proposed model can efficiently and effectively identify natural frequencies, damping ratios, and mode shapes of different engineering structures that belong to a population. Notably, the GNN-based model outperforms the traditional frequency domain decomposition method in terms of efficiency and surpasses a multilayer perceptron in accuracy, positioning it as a promising tool for population-based structural health monitoring.

 This research article was awarded Open Data for transparent practices. See the Data Availability Statement for details.

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.



1. Introduction

Structural health monitoring (SHM) has evolved into a useful tool for engineering practice, showcasing its importance through various applications in the management and maintenance of engineering structures (Farrar and Worden, 2012; An et al., 2019; Kamariotis et al., 2022). Although further steps are needed for standardization of SHM schemes, which are often bespoke to the structure at hand, the overall practice has matured, with the primary methods for extracting valuable information from SHM data categorized into physics-based and data-driven approaches (Avendaño-Valencia et al., 2017; Karakostas et al., 2024). The notable rise of data-driven approaches in recent years can be attributed not only to advancements in powerful deep learning (DL) tools, which favor computation, but also to the apparent robustness of schemes that capitalize on data against discrepancies between physical models and real-world structures (Haywood-Alexander et al., 2023; Cicirello, 2024). However, these approaches face challenges in interpretability and generalization, especially in civil engineering, where structures such as bridges are often uniquely designed to meet site-specific requirements (Lei et al., 2023). Training data-driven models with large datasets that represent a variety of structures can enhance their generalization capabilities. However, data from real structures is often lacking in quantity and diversity, as SHM has predominantly been applied to specific high-priority case studies due to budgetary and labor constraints. Consequently, the rich and diverse datasets necessary for training learning (data-driven) models are generally unavailable (Sun et al., 2020). To address this challenge, new sensing technologies, such as mobile sensing schemes (Jian et al., 2022, 2024; Stoura et al., 2023), are being developed to gather more measurement data from a greater number of structures. Additionally, advanced analysis methods can be used to more efficiently extract deeper insights from the available data. Treating structures as a population of shared features could enable learning and knowledge transfer among group members. For instance, although bridges are often unique, they can be categorized into broader typologies, such as truss or girder bridges. This approach is embodied in the recently introduced concept of population-based SHM (PBSHM) (Gosliga et al., 2022; Tsialiamanis et al., 2023, 2024; Bunce et al., 2024).

According to structural analysis theory, a population of structures can be represented by graphs that share certain morphological or topological characteristics, despite each instance comprising unique geometric or material properties. Following the logic of graph representations as suggested by Gosliga et al. (2021), graph neural networks (GNNs) can work as efficient representations for *PBSHM*. GNNs are a type of DL model that is specifically designed to operate on graph-structured data, comprising a set of nodes and edges that capture interrelations among elements of the network. Compared with traditional DL models, GNNs present several unique advantages (Sanchez-Lengeling et al., 2021) as follows: (1) *Flexibility*: The flexibility of the GNN architecture allows it to handle graphs with varying node counts and diverse connectivity patterns. This characteristic makes them particularly suitable for *PBSHM*-related datasets, where structures within a population share similarities but also possess individual differences. (2) *Accuracy*: GNNs are explicitly designed to handle data that is structured as graphs, which can process data in three different spatial levels: node level, edge level, and graph level. The message passing function of GNNs ensures more effective utilization of spatial information hidden in the data, consequently enhancing accuracy in processing graph-structured data. (3) *Interpretability*: We can naturally correlate data within a given dataset with the node-, edge-, and graph-level features of a GNN. This characteristic facilitates a more straightforward interpretation of both the model itself and the resulting outcomes.

Due to the above-mentioned advantages of GNNs, a number of studies have already attempted to integrate GNNs with forward analysis tasks. Examples include the use of GNNs to predict the shear stress in wall structures (Dupuy et al., 2023, 2024), airflow of wind energy systems (Mylonas, 2021; Duthé et al., 2023; Duthé et al., 2023), and estimating the main natural frequencies of truss structures given specific environmental conditions (temperature) and different member types (Tsialiamanis et al., 2022). For SHM, however, inverse problems play a more critical role, where the goal is to identify structural parameters from measurements of structures and, thus, evaluate the condition of a monitored structure (Gallet et al., 2022). To date, a paucity of studies involving GNNs for inverse problem solutions is noted. In this study,

we take this task on by introducing GNNs for a fundamental SHM task, namely Operation Modal Analysis (OMA), aiming to identify modal properties of a structure based solely on measured vibration response data (Reynders, 2012; Brincker and Ventura, 2015). As typically conducted within an OMA setting, the identified modal properties can be used for downstream tasks, such as damage identification (Hou and Xia, 2021; Gres et al., 2022; Greś et al., 2023) or response prediction (Lai et al., 2022). While conventional automated OMA methods have been shown to be effective for individual structures, they are typically applied independently and do not leverage population-level knowledge. The use of GNN for OMA could generalize across a population of structures, enabling rapid modal parameter inference with minimal additional computation. This makes it particularly advantageous for future big data applications, such as monitoring large-scale infrastructure networks.

Given the background introduced above, this study proposes a novel DL architecture, with the GNN serving as the pivotal building block, aimed at automatically performing output-only modal identification for structures within a population. The key contributions of this study include the below:

- We design the architecture of a GNN-based DL model, in which the structural population is modeled by a GNN. This model is trained to capture the meta behavior of the structural population, enabling it to output natural frequencies, damping ratios, and mode shapes from the input, which is the power spectral density (PSD) of the vibration acceleration on the structure nodes.
- We adopt the feature propagation (FP) algorithm to reconstruct the full-field acceleration PSD using partial measurements of a small subset of structure nodes, in an effort to reflect a realistic SHM context where only sparse dynamic response measurements are available. The proposed GNN-based model can then be used for modal identification based on the reconstructed PSD.
- We conduct a comprehensive set of numerical experiments using simulation data to evaluate the proposed GNN-based model. The study examines multiple GNN architectures, incorporates ablation studies, and evaluates performance under incomplete measurement conditions. Comparative analyses are carried out against established modal identification methods. Furthermore, we investigate the model's sensitivity to power spectrum resolution, training set size, and measurement noise, and assess its generalization capability across different structural populations. Both the accuracy and computational efficiency of the model are rigorously analyzed.

2. Methodology

2.1. Problem formulation

Modal identification of structural systems plays a fundamental role in the context of SHM. In general, the identification task aims at ascertaining the natural frequencies, damping ratios, and mode shapes of the monitored structure by processing measurements relating to dynamic response (structural output), such as accelerations, from a finite- and typically sparse set of degrees of freedom (DOFs). OMA schemes rely on the use of output-only information under the assumption of broadband and random unmeasured excitation. In keeping true to this requirement, this study employs Gaussian white noise to model the excitation sources of the simulated structures.

To further define the problem, let us denote the available acceleration measurements as $\mathbf{X}(t) \in \mathbb{R}^{N \times P}$, where t represents the time components, N denotes the number of monitored DOFs, and P denotes the number of available time samples per signal. In this study, we aim to learn a function $func(\cdot)$ that can identify natural frequencies, damping ratios, and mode shapes of the first k structural modes based on the measured signals, $\mathbf{X}(t)$. Then, the identification process can be expressed as:

$$[\hat{\mathbf{F}}, \hat{\mathbf{Z}}, \hat{\mathbf{\Phi}}] = func(\mathbf{X}(t)) \quad (2.1)$$

where $\hat{\mathbf{F}} \in \mathbb{R}_{>0}^{1 \times k}$, $\hat{\mathbf{Z}} \in \mathbb{R}_{>0}^{1 \times k}$, and $\hat{\mathbf{\Phi}} \in \mathbb{R}^{N \times k}$ denote identified natural frequencies (in Hz), damping ratios, and mode shapes, respectively.

Within an OMA context, P is usually chosen to reflect a large number of time samples, since typically long time series are required for the construction of the appropriate spectra. As a consequence, the input of the DL model, $\mathbf{X}(t)$, will comprise a high dimensionality, which is often undesirable as it compromises efficiency (Van Der Maaten et al., 2009). Therefore, in this study, we reduce the dimension of $\mathbf{X}(t)$ by employing the frequency domain representation of the measured signals, that is, the PSD. The converted PSD is denoted as $\mathbf{S}(f) \in \mathbb{R}_{\geq 0}^{N \times M}$, where f denotes the frequency components and M is the dimension of the PSD representation. Since the PSD contains only amplitude and not phase information of the time-history signals, only absolute mode shapes $|\hat{\Phi}| \in \mathbb{R}_{\geq 0}^{N \times k}$ can be identified from the proposed approach, which are nonetheless still meaningful for SHM tasks. Thus, the problem in this study can be reduced to learn a function $\text{func}(\cdot)$ that can be expressed as:

$$[\hat{\mathbf{F}}, \hat{\mathbf{Z}}, |\hat{\Phi}|] = \text{func}(\mathbf{S}(f)) \quad (2.2)$$

2.2. Model architecture

To obtain the modal identification function $\text{func}(\cdot)$ shown in Equation (2.2), we design a DL model to learn the mapping between $\mathbf{S}(f)$ and $[\hat{\mathbf{F}}, \hat{\mathbf{Z}}, |\hat{\Phi}|]$. The architecture of the proposed model is visualized in Figure 1.

As indicated in Figure 1, five DL blocks (marked in blue) are employed in the proposed model. Their description is as follows:

- *MLP 1*: The multilayer perceptron (MLP) 1 block serves as an encoder that compresses the input data \mathbf{S} (M -dimensional) into more compact hidden features $\mathbf{H1}$ (D -dimensional). This encoding process is widely adopted in the design of DL models (Goodfellow et al., 2016). The ablation study (Section 3.3) demonstrates that the usage of this block helps to improve the accuracy of the proposed model.
- *GNN*: The main building block in the proposed architecture is the GNN, which is constructed on the basis of engineering intuition. More specifically, as illustrated in Figure 2, we can reasonably designate the joint locations of a structure as the nodes of the GNN and the structural elements connecting those nodes (e.g., beams and truss bars) to correspond to the GNN edges. Unlike common DL models, such as MLPs or convolutional neural networks (GNNs), where the number of neurons is fixed, GNNs can process graphs of various configurations, corresponding to varying numbers of nodes and edges (Zhou et al., 2020), which is particularly beneficial for PBSHM tasks, since each structure usually comprises a different number of nodes and elements. The flexibility of GNNs comes from their message passing mechanism, which aggregates information from neighboring nodes rather than a fixed set of neurons, with the model parameters being shared across all nodes and edges. Once the skeleton of the GNN is built, the input quantities need to be determined on the basis of available measurements from the monitored systems. For the framework we propose, as




Figure 1. Architecture of the proposed model, in which the model input and output are marked in red, hidden features are marked in yellow, and deep learning blocks are marked in blue.




Figure 2. An example of the graph dataset used in this study. A truss structure can be naturally modeled as a graph. See the text about the GNN block for more details.

shown in Figure 2, this involves computation of the PSDs of vibration acceleration signals \mathbf{S} , as well as the absolute mode shapes $|\hat{\Phi}|$ as the node-level features of the GNN, while the natural frequencies $\hat{\mathbf{F}}$ and damping ratios $\hat{\mathbf{Z}}$ serve as the graph-level features of the GNN. This study adopts an encoder-decoder architecture, which implies that the GNN does not directly process \mathbf{S} as the input to produce $|\hat{\Phi}|$, $\hat{\mathbf{F}}$, and $\hat{\mathbf{Z}}$ as the output. Instead, the GNN takes the encoded features $\mathbf{H1}$ as the nodal input and executes two individual tasks, which in Figure 1 are correspondingly indicated as the graph-level task and the node-level task. The graph-level task corresponds to the “readout” operation, which is essential for graph-level downstream tasks (Gilmer et al., 2017), and serves for generating a single hidden feature $\mathbf{H2}$ by aggregating node features. In parallel, the node-level task performs the message passing operation, which is the GNN’s fundamental mechanism that allows nodes in a graph to exchange information with their neighbors in order to exploit the interrelations that lie latent in the graph-structured data (Wu et al., 2020). The product of the node-level task is the hidden feature $\mathbf{H3}$ on each node.

- *MLP 2–4:* Finally, three MLPs are employed as the decoder architecture in order to project the D -dimensional hidden features back to the physical domain, yielding the modal properties of the first k modes. The MLP 2 and MLP 3 blocks decode the graph-level hidden feature $\mathbf{H2}$, and thus output the graph-level modal properties, which are natural frequencies $\hat{\mathbf{F}}$ and damping ratios $\hat{\mathbf{Z}}$, respectively. The MLP 4 block decodes the node-level hidden feature $\mathbf{H3}$ and outputs the mode shapes $|\hat{\Phi}|$ of every node.

2.3. Loss function

To train our model, we employ an objective function based on the mean squared error MSE:

$$\mathcal{L} = \frac{\lambda_1}{kN} \sum_{i=1}^N \sum_{j=1}^k (|\Phi|_{ij} - |\hat{\Phi}|_{ij})^2 + \frac{\lambda_2}{k} \sum_{j=1}^k \left(\frac{\hat{\mathbf{F}}_j}{\mathbf{F}_j} - 1 \right)^2 + \frac{\lambda_3}{k} \sum_{j=1}^k \left(\frac{\hat{\mathbf{Z}}_j}{\mathbf{Z}_j} - 1 \right)^2 \tag{2.3}$$

where \mathbf{F} , \mathbf{Z} , and $|\Phi|$ are the target natural frequencies, damping ratios, and absolute mode shapes, respectively. These quantities are known during training, and the methods for deriving them will be elaborated in Section 2.4. λ_1 , λ_2 , and λ_3 are coefficients that are adopted to balance the contribution of the individual components related to \mathbf{F} , \mathbf{Z} , and $|\Phi|$ to the loss function. In this study, we determine λ_1 , λ_2 , and λ_3 by trial and error.

It is also noteworthy that, in this study, the PSDs (model input) and absolute mode shapes (model output) are normalized by scaling their amplitudes to a maximum of 1 unit before training to boost the generalization ability of the DL model. However, natural frequencies and damping ratios cannot be normalized by max normalization because their original values matter, and this can cause imbalance during model training. For instance, natural frequencies are generally much larger than damping ratios, so the loss term regarding natural frequencies will outweigh the damping ratio loss if they are not normalized. To diminish this possible imbalance, as shown in the loss function, we normalize natural frequencies and

damping ratios by dividing the estimated values (model output, which are $\hat{\mathbf{F}}$, $\hat{\mathbf{Z}}$) by their corresponding target values (\mathbf{F} , \mathbf{Z}).

2.4. Framework

Based on the abovementioned problem formulation, model architecture, and loss function, the proposed GNN-based *PBSHM* framework for modal identification is summarized in Figure 3, which shows that the framework is divided into two stages, namely model training and model implementation.

In the model training stage, a subset of the structural population is used as a training set. To establish such a population, numerical simulation is necessary. This is because, for real-world datasets, it remains impractical to monitor the dynamic response of all DOFs across many structures, whereas full-field measurements and mode shapes are required for supervised training. To address this challenge, we generate a synthetic dataset based on randomly sampled structural geometries and material parameters, modeled using the finite element method. The associated modal properties and complete acceleration PSDs are computed through numerical analysis. This simulation-derived dataset provides the basis for pretraining the proposed GNN model, allowing it to learn generalizable modal patterns across a diverse structural population and then perform OMA for unseen structures.

In practice, GNN models trained solely on idealized, noise-free data are expected to show reduced performance when applied to real-world OMA data due to the presence of measurement noise and modeling uncertainties. To address this domain gap, we adopt a “pretraining followed by fine-tuning” strategy. The model is first pretrained on large-scale synthetic data generated through eigenvalue analysis to capture general modal patterns across diverse structural configurations. It is then fine-tuned using a smaller set of manually validated OMA results from real structures, with finite element model updating employed to improve simulation fidelity. While this fine-tuning process cannot surpass the inherent error of the OMA labels, since the Bayes error bound (Fukunaga, 2013) imposes a theoretical lower limit, the knowledge gained during pretraining enables the model to generalize across new structures and potentially achieve comparable or improved performance to conventional OMA methods, particularly in scenarios such as sparse sensing or large-scale population monitoring. This trade-off is appropriate for a feasibility study whose aim is to demonstrate scalability and automation rather than deliver field-ready performance.

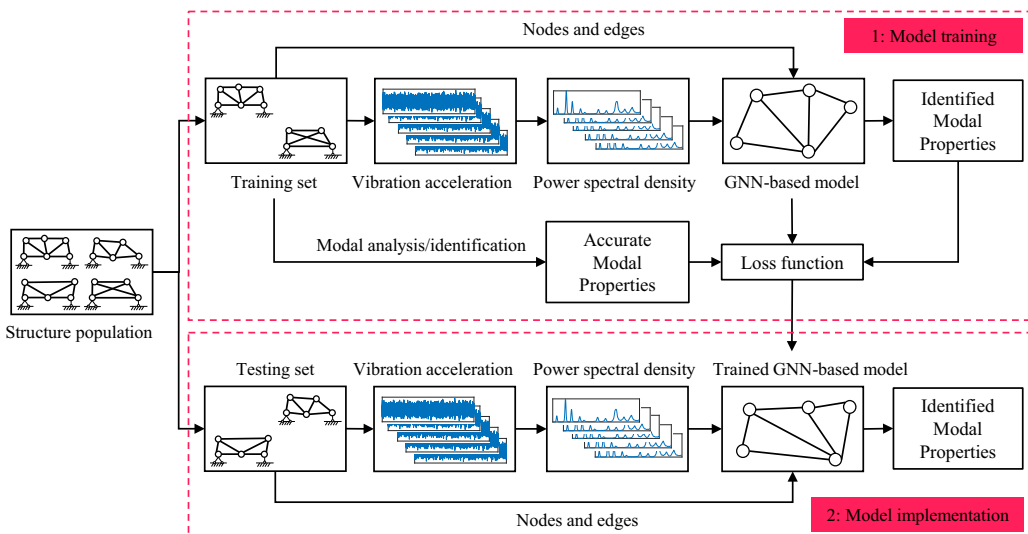


Figure 3. Framework for using the GNN-based model for population-based structural modal identification.

Due to the lack of a real-world PBSHM dataset, in this study, we only verify the proposed model with pure numerical experiments, so the fine-tuning part is not involved. More details about the dataset generation can be seen in [Section 3.1](#) below.

After the training is completed, the trained model can be readily used to automatically identify the modal properties of other unseen structures within the structural population, even though their configuration may be different from what is met in the training set, that is, corresponding to a different number of nodes and edges and connectivity. This identification process is end-to-end, automatic, and fast, delivering a main advantage over existing modal identification schemes. A more detailed comparison can be found in [Section 3.5](#).

3. Numerical validation

In this section, we validate the proposed GNN-based model using numerical simulation data, to address the following research questions:

- **Question 1:** Which type of GNN is most suitable for the population-based OMA task?
- **Question 2:** How do the DL components in the model architecture, particularly the GNN module, contribute to overall model performance?
- **Question 3:** Can the proposed model reliably perform modal identification when the measurement data are incomplete?
- **Question 4:** What are the strengths and limitations of the proposed model compared to existing modal identification methods?

3.1. Dataset description and implementation details

The extraction of datasets from real-world structural populations is nontrivial, although recent efforts in Asia (Moreu et al., 2018) and Europe (Limongelli et al., 2024) do involve denser instrumentation of structural populations (such as bridges). These endeavors are expected to deliver datasets that can support PBSHM tasks. However, these datasets have not been made public yet, so in this work, we initiate from simulated scenarios for two datasets that correspond to two different truss populations.

The first truss population consists of 2,600 trusses, assumed to be arranged within a trapezoidal boundary that is meant to approximate the geometrical configuration of simply-supported beam structures. Each truss is generated by randomly meshing the trapezoidal area with Delaunay triangles (Persson and Strang, 2004), which ensures that no point lies inside the circumcircle of any triangle, resulting in well-shaped elements with geometric stability. The first 2,500 trusses are used to train the model, among which 2,000 trusses are used as the training set and 500 trusses are used as the validation set. The last 100 trusses from the 2,600 trusses are used to test the model. The geometric boundary and some generated truss examples are shown in [Figure 4a,c](#), respectively. Based on the generated geometric configurations, the corresponding finite element models are straightforwardly created using truss elements with realistic structural parameters. The density and area of truss elements are constantly set as 8,015 kg/m³ and 0.5 m², respectively. In order to reflect varying material properties, the Young's modulus of the employed truss elements is set as a random number ranging from 100 to 300 GPa. [Figure 4a](#) further indicates the boundary conditions (simply-supported type) and the external excitation (Gaussian white noise) that are imposed on the bottom boundary. Linear time history analyses (Newmark- β method) are then performed to obtain the in-plane vertical nodal acceleration (response) time series, of 60 s duration, sampled at a time step of 0.005 s (200 Hz). Welch's method is next applied to convert the time series into PSDs, which serve as the input of the suggested DL architecture. The `pwelch` function in MATLAB is configured as follows: “`nfft`” = 2,048, “`noverlap`” = 1,024, hamming window applied, “`window length`” = 512, leading to the PSD having a resolution of 1,024 spectral lines. Moreover, in [Section 4.2](#), the proposed model is trained and tested under three different PSD resolutions, which are 1,024, 512, and 256, to investigate its sensitivity to the PSD resolution. An eigenvalue analysis is further conducted on all 2,600 finite element

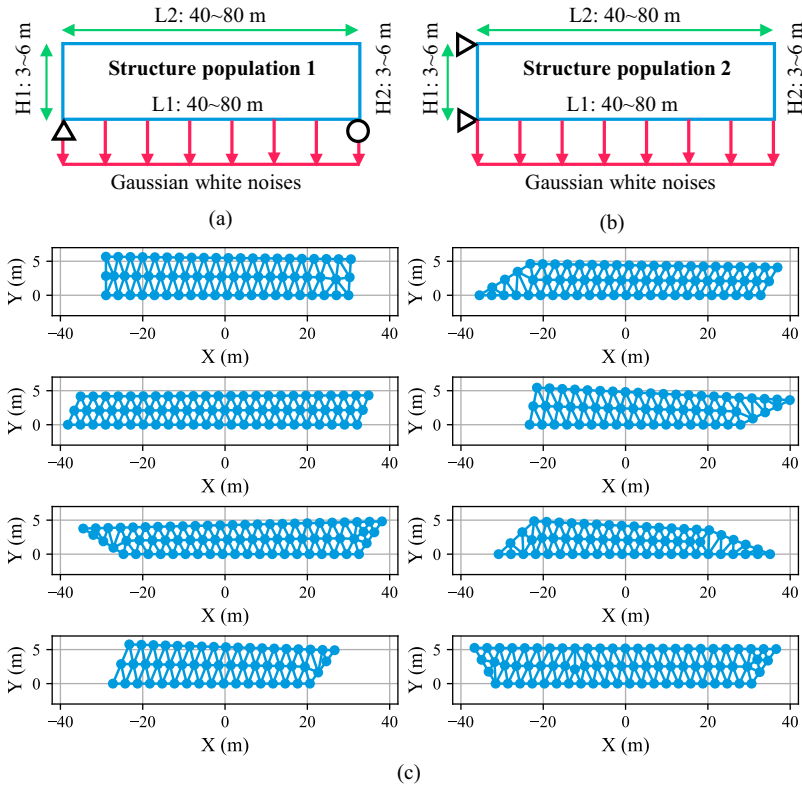


Figure 4. Visualization of the two generated datasets: (a) Geometric configuration meant to approximate a simply-supported truss population; (b) geometric configuration serving to approximate a cantilevered truss population; (c) some representative truss samples from the two datasets (only nodes and elements are displayed).

models, generating the reference modal properties of the primary (first four in our study) modes for model training and testing. The widely accepted Rayleigh damping model is employed to simulate damping effects. In particular, the modal damping ratio of the first and fifth vibration modes of every simulated truss is set to 0.01, which allows a subsequent calculation of the Rayleigh model coefficients (α, β). These two modes are chosen as they cover the modal range of interest for this study, which is the first four vibration modes. It should be noted that, for model training, the complete full-field nodal acceleration measurements are used, whereas the trained model is tested with both the complete and incomplete measurements. More details in this regard are offered in Section 3.4.

In order to further test the generalization ability of the proposed model, we simulate a separate structural population comprised of 100 samples that are meant to approximate the geometry of cantilevered trusses. This population is unseen for the learning PBSHM model, which is trained on the simply-supported structural population. As illustrated in Figure 4b, the second population only differs from the first population in the structural supports. Except for the structural supports, all the other settings, including geometric boundaries, material properties, and excitation, are the same as those of the first structural population. More details in this regard can be found in the following Section 4.4.

The environment we implement our model in is set up in PyTorch 2.0.1, CUDA 11.7, and DGL 1.1.1. The Adam optimizer is adopted to train our model with the default setting (learning rate: 0.002, first momentum decaying parameter: 0.9, second momentum decaying parameter: 0.999), and we use a NVIDIA GeForce RTX 3060 to accelerate the batch training process. The training batch size is 400, with

5,000 epochs deployed. The MLPs and GNNs used in our study all have three layers, and the adopted dimension is 64 for every layer. In the loss function, we set the coefficients λ_1 (regarding mode shapes), λ_2 (regarding natural frequencies), and λ_3 (regarding damping ratios) as 2, 1, and 1, respectively. More implementation details can be found in our GitHub repository (Jian, 2025), which will be made public upon publication of this work. It should be noted that this article focuses on demonstrating the feasibility of the proposed approach rather than presenting an optimal model with the best performance. We conducted some hyperparameter tuning, though nonexhaustive, as we aim to show functionality of the network without requiring such refined optimization. Interested readers may perform hyperparameter tuning using our publicly available code.

3.2. Effects of different GNN models

Depending on the type of message passing function that is adopted, different GNN modeling instances can be created. In this study, we compare three commonly used GNN models, namely graph convolutional networks (GCNs, proposed by Kipf and Welling, 2016), graph attention networks (GATs, proposed by Veličković et al., 2017), and GraphSAGE (proposed by Hamilton et al., 2017). To control the comparison, we only replace the GNN block in the model architecture shown in Figure 1 with different GNN models. Then, different models are trained using the K-fold cross-validation, which means we evenly divide the training set that consists of 2,500 trusses into 5 folds. In each training process, 4 folds of trusses (2,000 in number) are used to train the model, and the remaining 1 fold of trusses (500 in number) is used to validate the model. We train the model five times, so that all the five folds can serve as the validation set in turn. The mean values and standard deviation (SD) of the final validation loss and training time of the five training processes are shown in Figure 5.

Based on Figure 5, we observe that the GraphSAGE model demonstrates superior performance, as both its final validation loss and required training time are statistically smaller than the GCN and GAT models. This is because, according to Hamilton et al. (2017), GraphSAGE exhibits inductive learning capabilities, scalability via neighborhood sampling, and flexibility in aggregation methods. These features render it particularly suitable for large-scale and dynamic graphs, such as those employed herein within the PBSHM context, where efficiency and the ability to handle new nodes without re-training are crucial. GCNs are less scalable and inductive, while GATs offer sophisticated attention mechanisms at the cost of increased computational complexity. As a result, hereafter, we choose to adopt the GraphSAGE model as the GNN block, unless mentioned otherwise.




Figure 5. Comparison of different GNN models, using fivefold cross-validation during model training: (a) Final validation loss and (b) total training time.




Figure 6. Loss curves of one training process among the K -fold cross-validation: (a) Total training and validation loss; (b) different loss terms in the training loss; and (c) different loss terms in the validation loss.

Now that the particular type of GNN model is determined, it is necessary to elaborate on the details of the training and testing process in order to comprehensively report on the proposed model. Figure 6 shows the loss curves corresponding to one training process among the K -fold cross-validation. The total loss (Equation 2.3), the loss terms regarding absolute mode shapes ($|\hat{\Phi}|$), damping ratios (\hat{Z}), and natural frequencies (\hat{F}) are all displayed.

Figure 6 allows for the following observations:

- The training loss results are smaller than the validation loss in Figure 6a, which is reasonable as the validation loss is not used in backpropagation for updating the model weights.
- Nonetheless, the validation loss is introduced for deciding when to terminate the training process and for facilitating the comparison of different GNN models in this subsection, as well as different model architectures in the ablation study below. Although the training loss appears to continually decrease in Figure 6b, the validation loss appears to have converged in Figure 6c, which promoted termination of the training process after 5,000 epochs.

After the training is completed, we first use the trained model to identify the modal properties of truss configurations lying within the simply-supported structural population. Identification results of the testing set (100 trusses that do not occur in the training and validation set) are visualized in Figure 7, in which each dot represents the identification result per individual truss per mode, and different orders of modes are distinguished by different colors. Figure 7a shows the Modal Assurance Criterion (MAC) values of

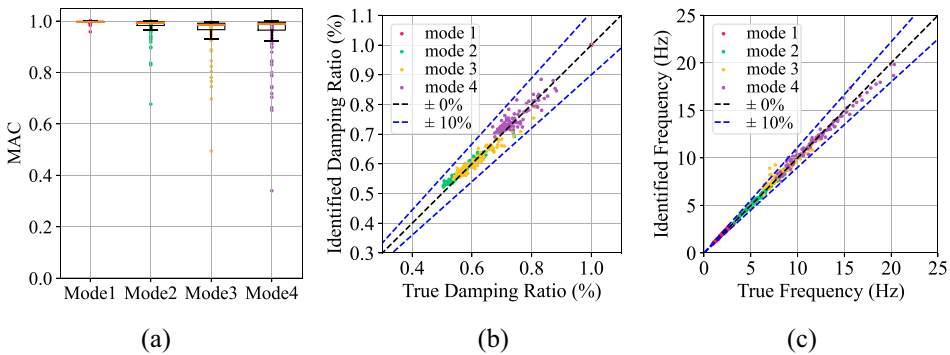


Figure 7. Performance of the trained model on the testing set of truss population 1 (simply supported): (a) Box plot of MAC values of identified mode shapes; (b) scatter plot of identified damping ratios; and (c) scatter plot of identified natural frequencies.

identified mode shapes, which is defined by Equation (3.1), along with their box plots. The closer the MAC value is to 1, the higher the accuracy of the mode shape identification.

$$\text{MAC}(|\hat{\phi}|, |\phi|) = \frac{(|\hat{\phi}|^T |\phi|)^2}{(|\hat{\phi}|^T |\hat{\phi}|)(|\phi|^T |\phi|)} \quad (3.1)$$

where $|\hat{\phi}|$ and $|\phi|$ are vectors of the identified and target absolute mode shapes, respectively.

The identified damping ratios and natural frequencies are shown in Figure 7b,c, respectively. In these figures, the further a point lies from the $\pm 0\%$ -error baseline, the larger the weighing error. Two extra reference lines representing the $\pm 10\%$ relative errors are also plotted in the figure to further evaluate the identification accuracy and reliability.

Figure 7 allows for the following findings:

- The training process is deemed as successful, because the trained model generalizes well on the testing set that consists of unseen trusses. Despite several instances of low identification accuracy, most identified mode shapes demonstrate MAC values close to 1, as shown in Figure 7a. In addition, most points plotted in Figure 7b,c are scattered near the $\pm 0\%$ -error baseline. Only a few cases show relative errors larger than 10%. This testing result is satisfactory for a purely data-driven approach, particularly considering that the trained model can generalize to unseen structures with different topologies and mechanical properties, and the identification process requires no human intervention.
- To quantitatively understand the performance of the model on the testing set, Table 1 presents the statistics of the MAC values of identified mode shapes and the relative errors (%) for the identified damping ratios and natural frequencies. As shown, the average MAC values for the identified mode shapes for all four modes are above 0.95, and the average errors corresponding to the damping ratios and natural frequencies are merely 1–2%. This indicates that the trained model statistically performs well in the modal identification of unseen structures among the same structural population.
- Table 1 also implies that the GNN-based model performs slightly worse when identifying modal parameters for higher-order modes. This is expected, as higher-order modes are inherently more challenging to identify. Another interesting observation is the remarkably high accuracy in identifying the damping ratios of the first mode (Mode 1). This is reasonable, given that the damping ratio of Mode 1 for every simulated truss in this study is consistently set at 0.01. With no variation in these values, the DL model can easily learn to identify them with exceptional accuracy.

3.3. Ablation study

To justify the model architecture shown in Figure 1, we perform an ablation study by comparing the following model variants:

- **No encoder:** Remove the encoder (MLP1) and retain the remaining blocks in the model architecture.

Table 1. Performance indicators on the testing set from the simply-supported structural population

	MAC of $ \hat{\Phi} $			Errors (%) of $\hat{\mathbf{Z}}$			Errors (%) of $\hat{\mathbf{F}}$		
	Mean	SD	Min	Mean	SD	Max	Mean	SD	Max
Mode 1	0.997	0.005	0.958	0.000	0.000	0.001	1.380	2.949	7.648
Mode 2	0.979	0.042	0.677	1.874	2.124	7.141	0.045	2.521	10.024
Mode 3	0.959	0.073	0.495	-0.640	2.585	7.196	1.710	5.015	26.979
Mode 4	0.949	0.098	0.340	1.248	2.844	8.642	0.039	3.879	11.497

Note. $\hat{\Phi}/\hat{\mathbf{Z}}/\hat{\mathbf{F}}$, identified mode shapes/damping ratios/natural frequencies; SD, standard deviation.

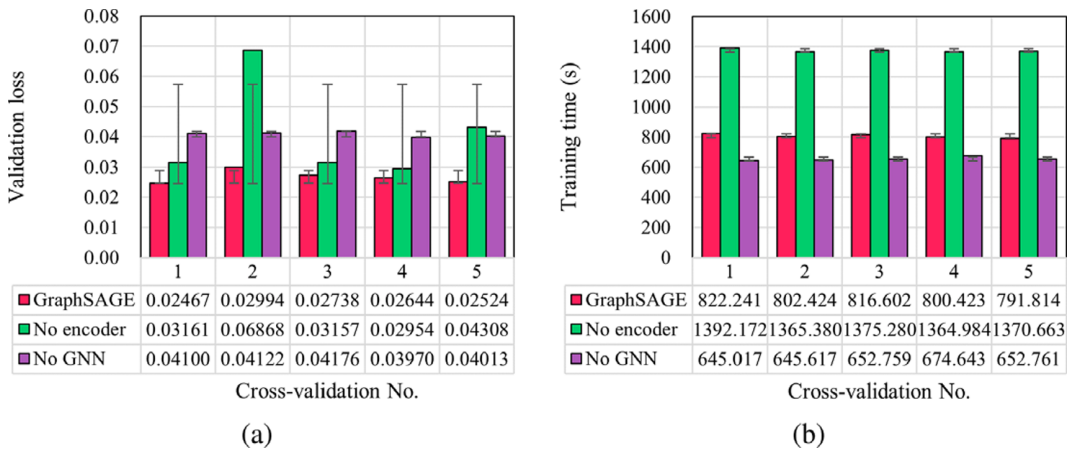


Figure 8. Comparison of different model architectures, using fivefold cross-validation during model training: (a) Final validation loss and (b) total training time.

- No GNN for mode shape identification:** Replace the GNN block with an MLP when converting the hidden features $\mathbf{H1}_{N \times D}$ to $\mathbf{H3}_{N \times D}$ in the model architecture. It should be noted that the GNN cannot be completely removed from the architecture because the identification of damping ratios and natural frequencies requires the graph readout operation. Without the GNN, the MLP that has fixed input dimensions will not be able to convert the hidden features $\mathbf{H1}_{N \times D}$ to $\mathbf{H2}_{1 \times D}$, making the architecture inapplicable to structures that possess a different number of nodes among a structural population.

To achieve this comparison, we train the originally proposed model and the two variants five times with the K-fold cross-validation method. The mean values and SD of the final validation loss and training time of the five training processes are shown in Figure 8.

Based on Figure 8, we can draw the following conclusions:

- The final validation loss indicates that the original model, whose architecture is illustrated in Figure 1, performs best among the tested variants because it exhibits the lowest validation loss, which justifies the design of the proposed model. The validation loss for the case of “No GNN” model is not only higher than the other two models in this ablation study but also higher than GCN and GAT models shown in Figure 5, confirming that the GNN does have an advantage in processing graph-structured data over the MLP.
- Removing the encoder and the GNN from the architecture increases the training time. Without the encoder to compress the model inputs, the GNN’s dimension must match the node PSD, making the model more complex and slower to train. Moreover, the absence of the encoder results in higher validation loss, indicating poorer model performance. Therefore, the original architecture is preferable.

3.4. Effects of incomplete measurements

In practice, the DOFs of civil engineering structures usually significantly outnumber the measurement points, so it is essential to investigate whether the proposed model can identify complete mode shapes from incomplete measurements. As stated previously, we construct the GNN corresponding to the analytical model of a structure. For those structural nodes that have no measurements, their node features are unknown, and the GNN model used in this study cannot work with unknown node features. To make the proposed model work in the case of unknown node features, we introduce the FP algorithm (Rossi et al., 2022) to fill the unknown features, which are nodal acceleration PSD, on the basis of known features and the graph structure. Applying FP transforms the incomplete measurements into complete ones, enabling our GNN-based model to conduct modal identification.

To evaluate the “FP + GNN” approach under reduced sensor availability, we modify the test dataset from truss population 1 by evenly removing PSD data from 82% of the nodes in each truss. The reason for evenly selecting nodes for removal, rather than randomly, is that sensors are typically evenly distributed on monitored structures in practice. The FP algorithm is then applied to the remaining 18% of nodes to estimate and fill the “unknown” features for the omitted 82% of nodes, generating filled, complete test datasets. The model, trained on the original complete training dataset (comprising 2,500 trusses from population 1), is subsequently tested on the filled complete test dataset (containing 100 trusses from population 1). **Figure 9** compares the mode shape identification results for a truss example from the original complete test dataset and the filled complete test dataset.

For further evaluation, the top two sections of **Table 2** offers statistics of the MAC values of GNN-identified mode shapes and the relative errors (%) of identified damping ratios and natural frequencies, under 82% unknown node features. By comparing the statistics in **Table 2**, we can conclude that:

- When there are 82% node features unknown, the identification accuracy drops, which is naturally expected. Moreover, the accuracy of identifying high-order mode shapes drops more than low-order mode shapes due to the greater complexity of high-order mode shapes.
- Compared with the identification of mode shapes, the identification of natural frequencies and damping ratios is much less influenced by the increase of unknown node measurements. This is reasonable because, in theory, natural frequencies and damping ratios can be successfully identified

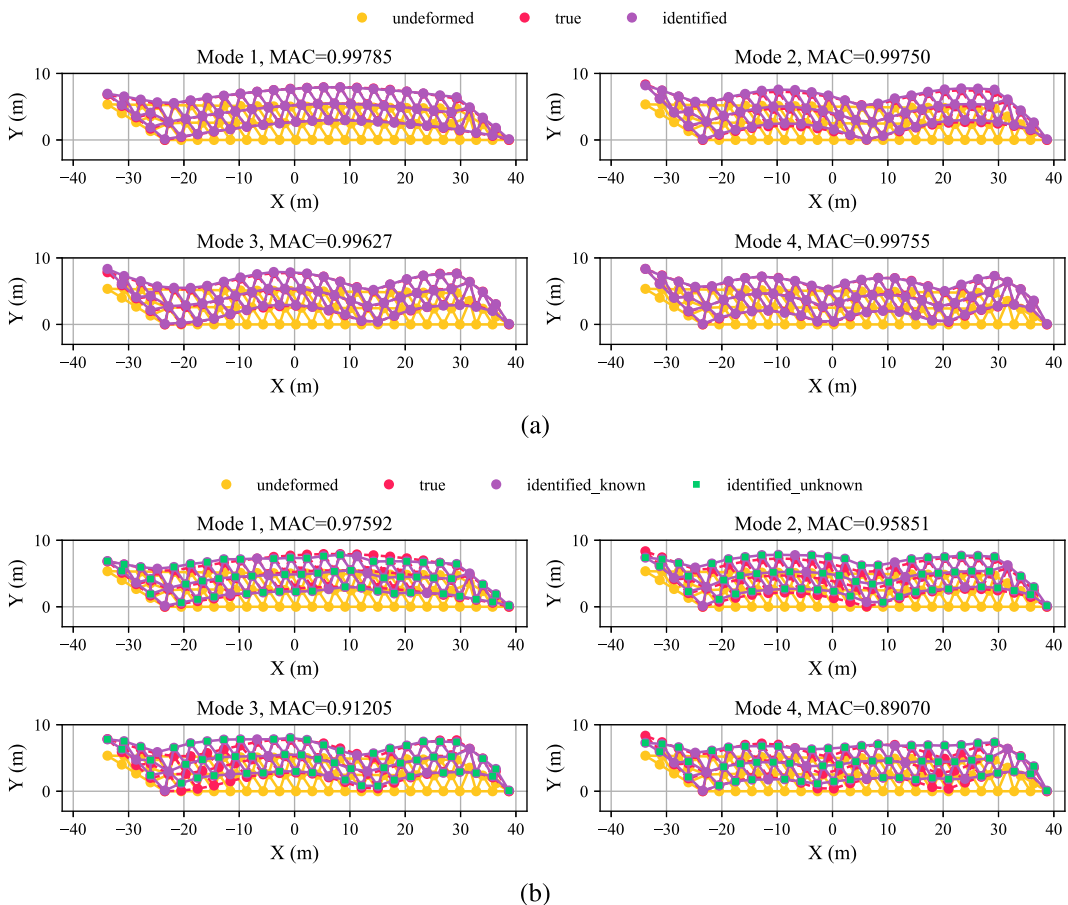


Figure 9. Mode shape identification results of one truss example with incomplete measurements: (a) 0% node features are unknown and (b) 82% node features are unknown.

Table 2. Performance indicators on the testing set from the simply-supported structural population

	MAC of $ \hat{\Phi} $			MAE of $ \hat{\Phi} $			Errors (%) of \hat{Z}			Errors (%) of \hat{F}		
	Mean	SD	Min	Mean	SD	Max	Mean	SD	Max	Mean	SD	Max
GNN-based model + 0% unknown node PSD												
Mode 1	0.997	0.005	0.958	0.031	0.017	0.132	0.000	0.000	0.001	1.380	2.949	7.648
Mode 2	0.979	0.042	0.677	0.061	0.035	0.238	1.874	2.124	7.141	0.045	2.521	10.024
Mode 3	0.959	0.073	0.495	0.077	0.048	0.292	-0.640	2.585	7.196	1.710	5.015	26.979
Mode 4	0.949	0.098	0.340	0.086	0.063	0.344	1.248	2.844	8.642	0.039	3.879	11.497
GNN-based model + 82% unknown node PSD												
Mode 1	0.967	0.019	0.876	0.106	0.036	0.217	-0.245	0.030	0.300	0.528	7.918	28.884
Mode 2	0.893	0.088	0.483	0.149	0.049	0.283	0.754	5.094	12.861	-2.551	5.852	24.777
Mode 3	0.834	0.105	0.354	0.183	0.067	0.453	0.292	6.092	17.982	-1.945	6.440	22.572
Mode 4	0.816	0.151	0.138	0.181	0.082	0.502	2.033	5.409	16.238	-2.723	7.902	23.424
EFDD + 0% unknown node acceleration												
Mode 1	1.000	0.000	1.000	0.000	0.000	0.003	206.770	105.065	413.323	-0.839	1.225	1.829
Mode 2	1.000	0.000	0.999	0.003	0.002	0.010	105.439	85.972	294.845	-0.519	0.423	0.434
Mode 3	0.989	0.085	0.156	0.014	0.028	0.269	55.312	51.222	200.788	-0.532	2.083	19.754
Mode 4	0.930	0.245	0.000	0.029	0.080	0.380	11.773	55.302	338.909	-1.713	10.125	59.839
EFDD + 82% unknown node acceleration												
Mode 1	0.982	0.025	0.852	0.055	0.026	0.165	209.714	105.521	393.802	-0.839	1.225	1.829
Mode 2	0.877	0.063	0.689	0.161	0.050	0.317	89.443	87.765	257.719	-0.519	0.423	0.434
Mode 3	0.781	0.110	0.492	0.204	0.074	0.417	56.103	52.826	211.309	-0.537	2.084	19.754
Mode 4	0.674	0.118	0.288	0.256	0.084	0.491	16.021	60.612	336.612	-1.661	10.565	59.839

Note. $\hat{\Phi}/\hat{Z}/\hat{F}$, identified mode shapes/damping ratios/natural frequencies; SD, standard deviation.

with only a few node measurements, whereas accurate identification of mode shapes requires measurements on more structural nodes. This observation implies that the trained model successfully learns the mechanism of modal identification.

- The modal identification accuracy is statistically acceptable despite being as high as 82% unknown node features, which manifests the robustness of the proposed model.

Finally, one might question whether a dataset with incomplete measurements should be used to train the proposed model to evaluate its learning and generalization abilities. However, this is unnecessary. As explained in [Section 2.4](#), numerical simulation of a structural population is essential for training the proposed model. The measurements, which are required as input for the use of this scheme on real-world structures, can be readily obtained as a result of a dynamic analysis under assumed ambient loads.

3.5. Comparison against existing modal identification methods

To thoroughly evaluate our model with respect to modal identification, we compare it against one of the most widely adopted modal identification methods, namely the enhanced frequency domain decomposition (EFDD) method (Brincker et al., 2000, 2001), which also operates in the frequency domain. The original EFDD method requires a manual peak-picking process to determine the modes that will be identified. For the population-based modal identification, however, manually picking peaks would be too time-consuming given the high number of trusses in our testing set. To tackle this problem, we adopt the automated EFDD algorithm proposed by Cheynet et al. (2017) to automatically identify mode shapes, damping ratios, and natural frequencies of the first four vibration modes for every simulated truss. The MATLAB codes of this automated EFDD algorithm will be made public in a dedicated GitHub repository (Jian, 2025) once this work is published. The bottom two sections of [Table 2](#) outline the statistical indicators of EFDD's modal identification performance on the testing set of the simply-supported structural population. Particularly, the last section shows EFDD's performance with incomplete measurements where acceleration on 82% structural nodes is evenly removed. In this scenario, mode shapes are initially identified using the 18% known node acceleration, followed by natural cubic spline interpolation and extrapolation to estimate the mode shapes for the remaining 82% of nodes with unknown acceleration data. We did not integrate the FP algorithm with EFDD for two key reasons. First, FP is inherently designed within the GNN framework, relying on graph-based representations to propagate information, whereas EFDD does not utilize any graph structure in its modal identification procedure. Second, the EFDD operates directly on raw time series data, which, in this study, comprises high-dimensional signals ($60 \text{ s} \times 200 \text{ Hz} = 12,000$ data points per node). Applying FP to such high-dimensional inputs risks introducing substantial imputation errors, which could ultimately degrade the performance of EFDD.

The comparison between the modal identification results of the GNN-based model and EFDD can provide the following findings:

- When structural dynamic response measurements are available at all nodes, representing an idealized (and not realistic) full-field measurement scenario, the GNN-based model achieves performance comparable to the EFDD in identifying mode shapes and natural frequencies. The EFDD scheme typically yields slightly higher MAC values and lower MAE for mode shapes, particularly in the first two modes. However, it is also more prone to outliers introduced by automated processing, leading to significantly larger maximum errors in some cases. In terms of natural frequency identification, both methods yield similar mean errors. The EFDD exhibits lower variability for modes 1–3 but higher SD and maximum error in mode 4. Notably, the GNN-based model significantly outperforms the EFDD in damping ratio estimation, as EFDD is known to be less reliable in this regard and produces large errors even under ideal conditions (Hasan et al., 2018).
- In the sparse sensing scenario, where only 18% of node features are available, the GNN-based model combined with FP demonstrates advantages over EFDD, supported by natural cubic spline interpolation. This is especially evident in higher-order mode shape identification, where the GNN

maintains higher MAC values and lower MAE. While the EFDD generally provides higher mean accuracy and lower variability in natural frequency estimation under sparse data conditions, it still exhibits substantial variability in mode 4. Furthermore, EFDD's damping ratio estimates remain highly inaccurate, whereas the GNN-based model delivers stable and reliable results. These findings underscore the GNN model's capability to effectively infer missing information by exploiting the spatial correlations embedded in the graph-structured data.

- In terms of the processing time, our GNN-based model is much more efficient when performing automated modal identification for a structural population. Part of the computational gain lies in the fact that the automated EFDD requires substantial computational effort to construct the PSD matrix, perform singular value decomposition, and automatically identify peaks. In contrast, the GNN-based model identifies modal parameters through a single forward pass, streamlining the entire process. For example, in this numerical experiment, we run both the automated EFDD and our GNN-based model on the same Central Processing Unit (CPU) under identical settings to identify the modal properties of 100 trusses sequentially. The automated EFDD spends 640.235 s on the dataset with 0% unknown node features and 39.947 s on the dataset with 82% unknown node features. In contrast, our GNN-based model completes the same task in just about 2.5 s for both datasets.
- If we count in the training time of the GNN-based model, which is on average 807.185 s as shown in [Figure 5b](#), our model requires more time than FDD. Considering the training time would, however, be an unfair comparison, since the intention (and main benefit) is to deploy the GNN-based modal identification approach in testing mode. This advantage becomes even more pronounced when dealing with larger structural populations, rendering the proposed GNN-based model an efficient tool for population-based modal identification.
- In conclusion, both the GNN-based model and the conventional EFDD method exhibit strong capabilities in modal parameter identification, each with distinct strengths and limitations. While neither approach consistently outperforms the other across all metrics, the GNN-based model emerges as a robust and reliable alternative, particularly under conditions of sparse sensor coverage, where its performance advantage becomes more evident.

4. Sensitivity analysis

In this section, we perform a sensitivity analysis of the proposed GNN-based model under various input configurations to address the following research questions:

- **Question 1:** Can the proposed model remain robust under realistic challenges, such as low PSD resolution, measurement noise, and a limited number of training structures?
- **Question 2:** Does the proposed model exhibit sufficient generalization capability across different structural populations?

4.1. Effect of training set size

As mentioned above, the proposed population-based modal identification method necessitates data from a group of structures to train the GNN-based model. Intuitively, a larger training set provides more diversified data, which in turn enhances the performance of the trained model. While this is easier to configure when employing a simulated dataset, where populations are artificially generated, the situation is different when aiming to use real-world datasets for training. In real-world applications, the number of monitored structures that can serve for model training is typically scarce, while the number of available sensors is also limited. The latter challenge is often tackled via the use of mobile or wide-coverage (such as computer vision measurements) sensor setups. In accounting for the first challenge, we here investigate the impact of the training set size on the model's performance. In this section, we use the first 20, 200, and 500 trusses (1, 10, and 25% of the original training set size) among the simply-supported population to train the model. Accordingly, the training batch size is changed to 4, 40, and 100 to accommodate the

shrinkage of the training set. Although these subsets of trusses are not particularly selected, they can still represent the entire truss population because the randomness in generating the dataset follows a uniform distribution, meaning the probability of each sample is equal.

The trained models are then used to identify the modal properties of 100 trusses from the testing set, with the results shown in Table 3.

By comparing the results of different training datasets shown in Table 3, we find that:

- As anticipated, reducing the size of the training set will decrease the accuracy of modal identification. However, even with a training set that amounts to only 20 trusses (1% of the original size), the decrease in accuracy can be deemed barely acceptable. A training set that amounts to 500 trusses (20% of the original size) can produce results slightly worse than those obtained with the 2,500-truss dataset. Thus, while a larger training set improves modal identification, a smaller one can still be effective when more data are unavailable.
- Compared against mode shapes, the decrease in the size of the training set induces higher errors in the identification of damping ratios and natural frequencies, especially in terms of maximum relative errors. A possible explanation for this is postulated in that mode shapes form spatial quantities, whose nature is more graph-structured than natural frequencies and damping ratios (which are global variables). Thus, the GNN can likely more robustly tackle mode shape identification, even at smaller training sets.

4.2. Effects of PSD resolution

The resolution of acceleration PSD may influence the GNN-based model's performance, especially for trusses with closely spaced modes. To evaluate this effect, we trained and tested the GNN-based model

Table 3. Performance indicators on the testing set from the simply-supported structural population—different sizes of training datasets

	MAC of $ \hat{\Phi} $			Errors (%) of $\hat{\mathbf{Z}}$			Errors (%) of $\hat{\mathbf{F}}$		
	Mean	SD	Min	Mean	SD	Max	Mean	SD	Max
Training with 2,000 trusses									
Mode 1	0.997	0.005	0.958	0.000	0.000	0.001	1.380	2.949	7.648
Mode 2	0.979	0.042	0.677	1.874	2.124	7.141	0.045	2.521	10.024
Mode 3	0.959	0.073	0.495	-0.640	2.585	7.196	1.710	5.015	26.979
Mode 4	0.949	0.098	0.340	1.248	2.844	8.642	0.039	3.879	11.497
Training with 500 trusses									
Mode 1	0.995	0.008	0.951	0.207	0.066	0.406	0.838	7.668	56.527
Mode 2	0.967	0.058	0.645	0.941	3.576	12.637	-0.034	7.730	54.447
Mode 3	0.940	0.094	0.456	2.422	4.637	14.455	0.320	6.061	20.738
Mode 4	0.912	0.143	0.320	0.549	4.013	13.719	-1.650	8.080	27.368
Training with 200 trusses									
Mode 1	0.992	0.013	0.915	-0.343	0.021	0.416	2.905	13.784	124.552
Mode 2	0.954	0.072	0.600	3.272	4.863	14.164	0.944	11.440	95.898
Mode 3	0.901	0.117	0.468	1.660	6.573	16.722	1.275	8.938	51.306
Mode 4	0.902	0.143	0.370	-1.174	5.399	17.342	0.605	11.527	82.717
Training with 20 trusses									
Mode 1	0.967	0.030	0.867	0.723	0.223	1.072	14.045	20.823	72.099
Mode 2	0.893	0.114	0.421	3.759	6.131	20.756	9.486	20.574	72.704
Mode 3	0.799	0.117	0.330	0.310	7.798	22.153	2.722	13.763	35.252
Mode 4	0.809	0.161	0.321	0.557	5.448	13.336	8.814	21.963	69.904

Note. $\hat{\Phi}/\hat{\mathbf{Z}}/\hat{\mathbf{F}}$, identified mode shapes/damping ratios/natural frequencies; SD, standard deviation.

Table 4. Performance indicators on the testing set from the simply-supported structural population—different PSD resolutions

	MAC of $ \hat{\Phi} $			Errors (%) of $\hat{\mathbf{Z}}$			Errors (%) of $\hat{\mathbf{F}}$		
	Mean	SD	Min	Mean	SD	Max	Mean	SD	Max
PSD resolution: 1,024									
Mode 1	0.997	0.005	0.958	0.000	0.000	0.001	1.380	2.949	7.648
Mode 2	0.979	0.042	0.677	1.874	2.124	7.141	0.045	2.521	10.024
Mode 3	0.959	0.073	0.495	−0.640	2.585	7.196	1.710	5.015	26.979
Mode 4	0.949	0.098	0.340	1.248	2.844	8.642	0.039	3.879	11.497
PSD resolution: 512									
Mode 1	0.997	0.003	0.974	−1.752	0.051	1.874	0.587	3.993	17.186
Mode 2	0.984	0.026	0.846	0.870	2.621	14.321	−0.053	3.509	10.138
Mode 3	0.948	0.108	0.264	−0.038	3.372	18.060	−0.323	4.818	24.537
Mode 4	0.944	0.104	0.461	1.061	3.065	11.024	−0.551	3.611	17.453
PSD resolution: 256									
Mode 1	0.997	0.004	0.962	−3.089	1.696	19.816	−0.500	3.671	14.912
Mode 2	0.981	0.033	0.798	0.785	3.626	26.181	−1.529	3.522	13.984
Mode 3	0.959	0.073	0.476	−1.425	3.122	13.355	−0.578	5.150	29.304
Mode 4	0.940	0.104	0.537	0.923	4.333	32.213	−0.497	5.589	27.244

Note. $\hat{\Phi}/\hat{\mathbf{Z}}/\hat{\mathbf{F}}$, identified mode shapes/damping ratios/natural frequencies; SD, standard deviation.

using datasets with PSD resolutions of 1,024, 512, and 256, respectively. The corresponding testing results are presented in Table 4.

As shown in Table 4, reducing PSD resolution does impact the performance of the GNN-based model, although the effect is relatively minor. The most noticeable consequence is an increase in the maximum errors of the identified damping ratios and natural frequencies. This outcome is anticipated, as lower PSD resolutions lead to a reduced frequency resolution and broader spectral peaks, which diminish the accuracy of natural frequency and damping ratio identification. This effect is particularly pronounced for closely spaced modes, causing an increase in maximum identification errors while having an insignificant influence on the mean errors. Therefore, higher PSD resolutions are recommended for achieving greater accuracy in modal identification.

4.3. Effects of measurement noise

In real-world applications, noise corruption is inevitable when acquiring signals that measure the vibration acceleration of structures. To test our model's sensitivity to measurement noise, we artificially pollute the acceleration time series generated by the testing set by injecting white Gaussian noise of which the root mean square is 10% of the signal power. Then, we use the Welch's method to obtain the PSD of the polluted time series and input these to the model trained with noise-free data for modal identification. Table 5 reports the statistics of the MAC values of GNN-identified mode shapes and the relative errors (%) of identified damping ratios and natural frequencies on the noise-polluted testing set. The automated EFDD algorithm is also tested on the noise-polluted dataset for comparison.

Results in Table 5 demonstrate that:

- The introduction of measurement noise slightly reduces the accuracy of identifying mode shapes and damping ratios. However, the average and maximum errors for the identified natural frequencies increase more heavily. This discrepancy may arise because mode shapes and damping ratios are relative measures, while natural frequencies are absolute values. The addition of noise alters the

Table 5. Performance indicators on the testing set from the simply-supported structural population—adding 10% white Gaussian noise

	MAC of $ \hat{\Phi} $			MAE of $ \hat{\Phi} $			Errors (%) of \hat{Z}			Errors (%) of \hat{F}		
	Mean	SD	Min	Mean	SD	Max	Mean	SD	Max	Mean	SD	Max
GNN + 0% noise												
Mode 1	0.997	0.005	0.958	0.031	0.017	0.132	0.000	0.000	0.001	1.380	2.949	7.648
Mode 2	0.979	0.042	0.677	0.061	0.035	0.238	1.874	2.124	7.141	0.045	2.521	10.024
Mode 3	0.959	0.073	0.495	0.077	0.048	0.292	−0.640	2.585	7.196	1.710	5.015	26.979
Mode 4	0.949	0.098	0.340	0.086	0.063	0.344	1.248	2.844	8.642	0.039	3.879	11.497
GNN + 10% noise												
Mode 1	0.988	0.010	0.949	0.059	0.032	0.177	0.000	0.000	0.001	−0.927	10.424	36.918
Mode 2	0.954	0.050	0.755	0.127	0.085	0.353	4.089	4.672	14.550	−3.293	8.311	34.635
Mode 3	0.841	0.100	0.438	0.195	0.065	0.410	1.629	6.093	20.907	−4.365	8.956	33.147
Mode 4	0.866	0.133	0.444	0.189	0.141	0.508	1.773	4.160	12.774	−7.026	10.149	42.899
EFDD +10% noise												
Mode 1	0.999	0.002	0.976	0.016	0.007	0.037	199.054	109.329	441.823	−0.815	1.223	1.829
Mode 2	1.000	0.000	0.999	0.008	0.006	0.049	108.654	83.425	292.954	−0.519	0.415	0.434
Mode 3	0.988	0.085	0.155	0.019	0.028	0.271	50.755	50.897	146.841	−0.532	2.083	19.754
Mode 4	0.928	0.247	0.001	0.027	0.074	0.376	4.983	59.234	330.402	−1.482	10.352	59.839

Note. $\hat{\Phi}/\hat{Z}/\hat{F}$, identified mode shapes/damping ratios/natural frequencies; SD, standard deviation.

absolute amplitude of the resulting PSDs without perceptibly changing the relative amplitude between the PSDs of different nodes.

- For mode shape and natural frequency identification, the automated EFDD outperforms the GNN-based model on the noise-polluted dataset. This is due to EFDD’s use of singular value decomposition on the PSD matrix, which effectively mitigates the adverse effects of measurement noise. However, the damping ratio identification results from the automated EFDD still display significant errors.

4.4. Effects of different structural populations

Last but not least, we investigate the generalization capability of the proposed model on different structural populations. We apply the model trained on the simply-supported truss population, as shown in Figure 4a, to the cantilevered truss population, as shown in Figure 4b. As an example, Figure 10 illustrates the identified mode shapes of a truss from the cantilevered population. The modal identification results of all (100) trusses are visualized in Figure 11, where each dot represents the identification result for a certain mode of a truss. Table 6 statistically summarizes the identification results illustrated in Figure 11.

The identification results on the cantilevered structural population are clearly worse than those on the simply-supported structural population. This disparity is expected, given the differences between the cantilevered and simply-supported datasets, with the former not being utilized during the training of the GNN-based model. Nonetheless, the identification results are not entirely unsatisfactory. For instance,

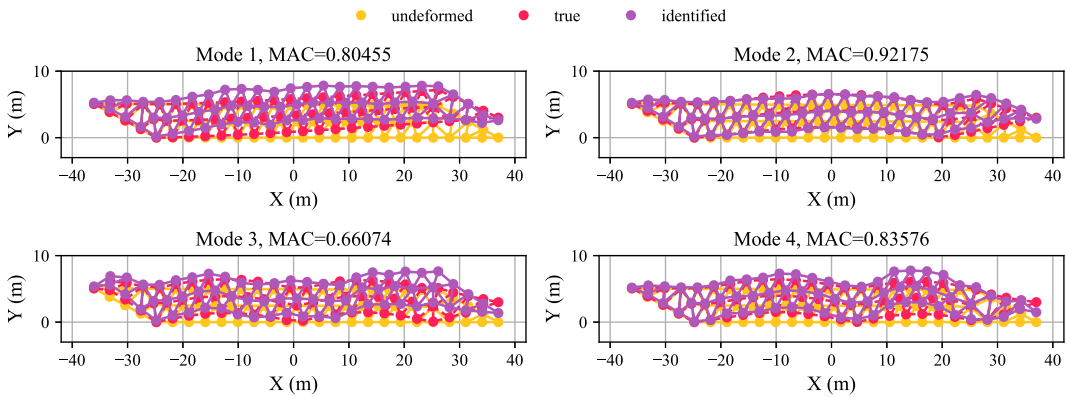


Figure 10. Mode shape identification results of a truss from the cantilevered population.

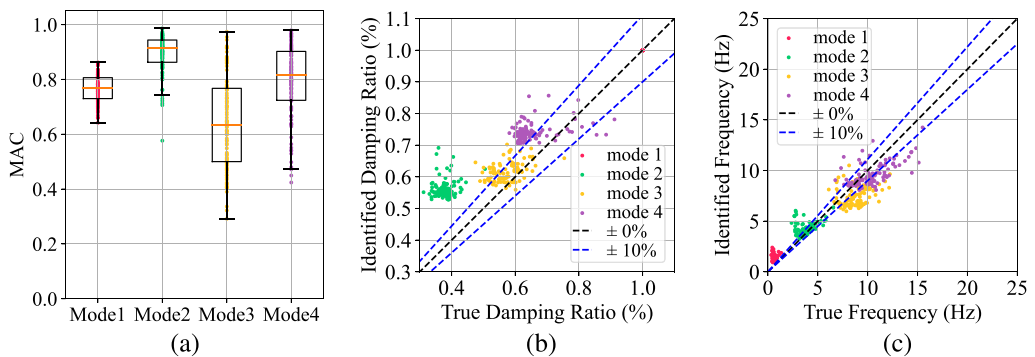


Figure 11. Performance of the trained model on the cantilevered truss dataset: (a) Box plot of MAC values of identified mode shapes; (b) scatter plot of identified damping ratios; (c) and scatter plot of identified natural frequencies.

Table 6. Performance indicators on the testing set from different truss populations

	MAC of $ \hat{\Phi} $			Errors (%) of \hat{Z}			Errors (%) of \hat{F}		
	Mean	SD	Min	Mean	SD	Max	Mean	SD	Max
Testing set from the simply-supported truss population									
Mode 1	0.997	0.005	0.958	0.000	0.000	0.001	1.380	2.949	7.648
Mode 2	0.979	0.042	0.677	1.874	2.124	7.141	0.045	2.521	10.024
Mode 3	0.959	0.073	0.495	-0.640	2.585	7.196	1.710	5.015	26.979
Mode 4	0.949	0.098	0.340	1.248	2.844	8.642	0.039	3.879	11.497
Testing set from the cantilevered truss population									
Mode 1	0.766	0.050	0.642	0.000	0.000	0.001	116.734	82.943	448.017
Mode 2	0.893	0.070	0.576	49.318	12.350	92.681	16.232	28.739	117.220
Mode 3	0.641	0.168	0.291	8.064	8.174	35.931	-13.978	14.690	37.741
Mode 4	0.794	0.136	0.424	14.082	9.506	32.865	-6.207	13.775	49.250

Note. $\hat{\Phi}/\hat{Z}/\hat{F}$, identified mode shapes/damping ratios/natural frequencies; SD, standard deviation.

although the relative errors of identified natural frequencies and damping ratios are high, as shown in Table 6, the identified values do not deviate significantly from their target values as Figure 11b,c shows. This indicates that the GNN-based model exhibits a degree of generalization capability, enabling it to leverage commonalities between two distinct yet similar structural populations.

To further explore how the similarity between these two truss populations affects the performance of the trained model, Figures 12–14 present probability histograms of the sum of squared mode shapes, the true natural frequencies, and damping ratios from the three datasets used in this study, respectively. The following observations can be made:

- As seen in Figures 12a, 13a, and 14a, the data distributions of the training and testing datasets from truss population 1 closely overlap. This similarity enables the trained model to generalize effectively to the testing dataset from the same population.

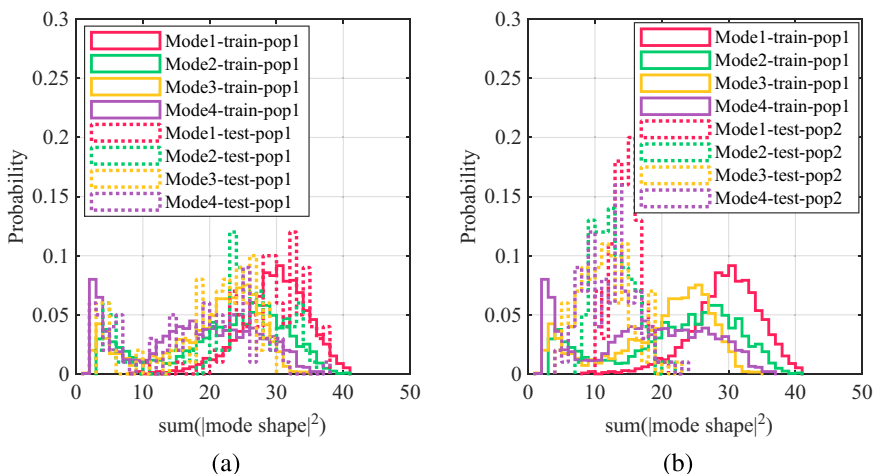


Figure 12. Probability histogram of the sum of squared mode shapes that belong to: (a) Training dataset of truss population 1 and testing dataset of truss population 1; (b) training dataset of truss population 1 and testing dataset of truss population 2.

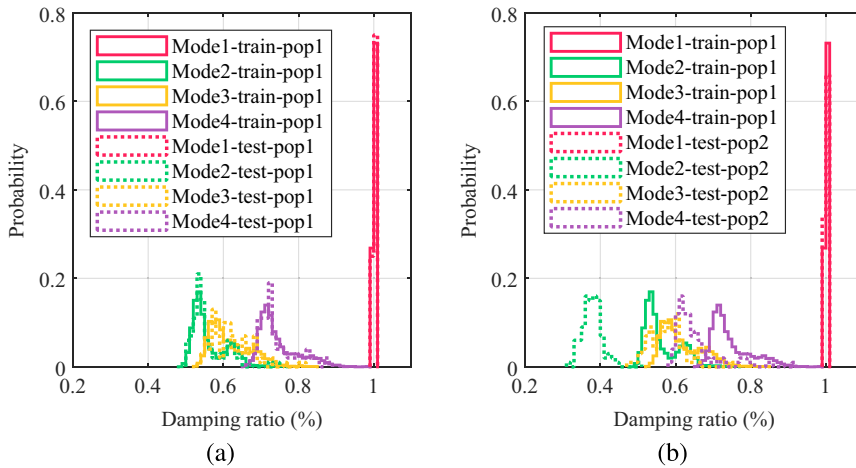


Figure 13. Probability histogram of true damping ratios that belong to: (a) Training dataset of truss population 1 and testing dataset of truss population 1; (b) training dataset of truss population 1 and testing dataset of truss population 2.

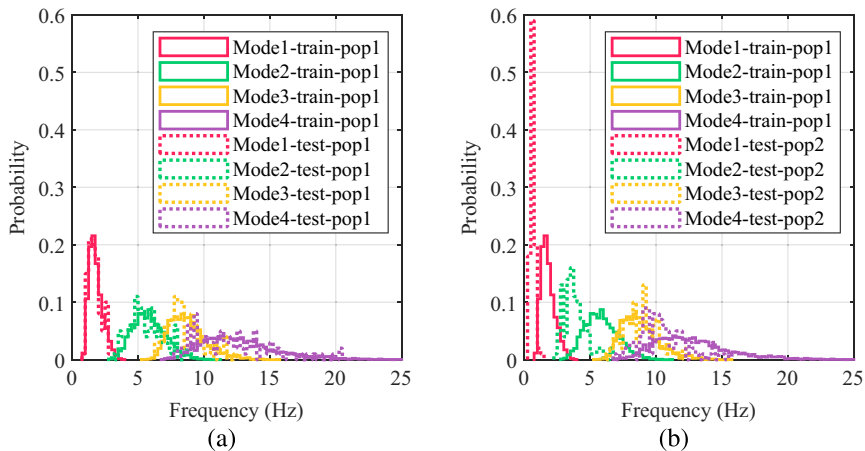


Figure 14. Probability histogram of true natural frequencies that belong to: (a) Training dataset of truss population 1 and testing dataset of truss population 1; (b) training dataset of truss population 1 and testing dataset of truss population 2.

- In contrast, Figures 12b, 13b, and 14b reveal significant differences between the data distributions of the testing dataset from truss population 2 and the training dataset from truss population 1. These discrepancies lead to poorer generalization of the model to the testing dataset of truss population 2. Furthermore, the extent of distributional deviation correlates with performance degradation. For instance, in Figure 13b, the damping ratio distribution of Mode 2 in truss population 2 has the largest deviation from that of truss population 1, resulting in the highest identification errors for Mode 2’s damping ratios, as shown in Table 6.

5. Conclusions

To exploit the information contained within a structural population for PBSHM, this study develops an automated operational modal analysis method capable of automatically identifying modal properties of

structures among a population by using a GNN-based DL model. A series of numerical experiments is carried out to validate the proposed method. The main findings emerging from this study are summarized as follows:

(1) *Model architecture.* The model comparison suggests that the GraphSAGE model outperforms the GCN and GAT models. The ablation study reveals that using an MLP to encode the PSD input of the model is beneficial. More importantly, the no-GNN model is inferior to all three GNN-based models, confirming the superiority of GNNs in processing graph-structured data.

(2) *Modal identification performance.* Trained on the modal properties of structures within a simulated structural population, the proposed GNN-based model effectively and efficiently identifies natural frequencies, damping ratios, and mode shapes of various structures within the same population, despite the variation of input PSD resolutions. When provided with a sufficiently large training set and tested on the same population of structures, the GNN-based model achieves accuracy comparable to the classic EFDD method while significantly surpassing it in computational efficiency. The performance is robust even in the presence of measurement noise and spatially sparse measurements. To tackle the issue of measurement sparsity, the FP algorithm is needed to reconstruct the full-field node input features, specifically the PSD of node acceleration, before using the proposed model to process incomplete measurements. Moreover, the “FP + GNN” approach is more accurate than the EFDD when 82% of node features are unknown. It should be noted, however, that the robustness and efficiency of the GNN-based model depend on effective training, so manual effort is required to prepare representative and diverse training data.

(3) *Limitations and future directions.* When training data are insufficient, or the testing structure belongs to a different population, the model’s accuracy declines, indicating limited generalization ability unless trained with sufficient and representative data. Notably, the model’s performance is strongly influenced by the degree of similarity between the data distributions of the training and testing sets. This is expected due to the purely data-driven nature of the proposed approach. Future research will focus on incorporating physical information to enhance the model’s generalization ability and evaluate uncertainty propagation. For example, according to the modal analysis theory, the vibration mode shapes of structures are orthogonal with respect to the mass matrix. Such properties could be incorporated into the loss function to train the model using a physics-informed approach. Another direction is to incorporate structural geometric properties directly into the FP process or use them as edge attributes in the GNN, which could further enhance mode shape identification and improve generalization across diverse structural populations. In the current study, such properties are not explicitly provided; instead, their effects are implicitly captured in the vibration responses (i.e., the PSDs) that form the basis for modal identification. A further limitation is that the model can only identify absolute mode shapes, instead of signed mode shapes, owing to the use of the auto-PSD that lacks phase information. While this is not prohibitive for SHM purposes, additional node features, such as phase angle spectra of node acceleration, or edge features like the cross PSD between the accelerations of two connected nodes, could be utilized to address this issue, as the sign of mode shapes is closely associated with the phase information in acceleration signals. It would further be valuable to validate the proposed method using real-world data, as noisy modal identification results from conventional OMA methods could affect the fine-tuning stage of the proposed model. This, however, does not form part of this preliminary study. Given the inherently incomplete nature of real-world data, the FP algorithm also requires thorough investigation to better handle sparse measurements in practice. Lastly, the proposed method admits refinements for reliable use within a damage detection context. The main aim of this work is to present an approach that can be used as a fast and reliable modal identification tool for extended bridge populations, inspired by the PBSHM paradigm. Future research could focus on integrating GNN with PBSHM to develop more effective approaches for damage detection.

Data availability statement. Demonstrative MATLAB and Python codes that implement the proposed method are openly available at <https://doi.org/10.5281/zenodo.16779978>.

Author contribution. Conceptualization: Eleni Chatzi and Xudong Jian. Methodology: Xudong Jian, Yutong Xia, Gregory Duthé, Kiran Bacsa, and Wei Liu. Data curation: Xudong Jian. Formal analysis: Xudong Jian. Software: Yutong Xia. Data visualization: Xudong Jian. Writing—original draft: Xudong Jian. Writing—review and editing: Eleni Chatzi, Xudong Jian, Kiran Bacsa, Yutong Xia, and Gregory Duthé. Supervision: Eleni Chatzi. Funding acquisition: Eleni Chatzi. Project administration: Eleni Chatzi. All authors approved the final submitted draft.

Funding statement. The research was conducted at the Singapore-ETH Centre, which was established collaboratively between ETH Zurich and the National Research Foundation Singapore. This research is supported by the National Research Foundation, Prime Minister's Office, Singapore, under its Campus for Research Excellence and Technological Enterprise (CREATE) program.

Competing interests. The authors declare none.

Ethical standard. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

References

- An Y, Chatzi E, Sim S-H, Laflamme S, Blachowski B and Ou J (2019) Recent progress and future trends on damage identification methods for bridge structures. *Structural Control and Health Monitoring* 26 (10), e2416.
- Avendaño-Valencia LD, Chatzi EN, Koo KY and Brownjohn JM (2017) Gaussian process time-series models for structures under operational variability. *Frontiers in Built Environment* 3, 69.
- Brincker R and Ventura CE (2015) Front matter. In *Introduction to Operational Modal Analysis* (pp. i–xii). John Wiley & Sons, Ltd. <https://doi.org/10.1002/9781118535141.fmatter>.
- Brincker R, Ventura CE and Andersen P (2001) Damping estimation by frequency domain decomposition. *Proceedings of IMAC 19: A conference on structural dynamics: February 5–8, 2001, Hyatt Orlando, Kissimmee, Florida, 2001*, 698–703.
- Brincker R, Zhang L and Andersen P (2000, 2000) *Modal Identification from Ambient Responses Using Frequency Domain Decomposition IMAC 18: Proceedings of the International Modal Analysis Conference (IMAC)*. San Antonio, Texas, USA February 7–10, pp. 625–630.
- Bunce A, Brennan DS, Ferguson A, O'Higgins C, Taylor S, Cross EJ, Worden K, Brownjohn J and Hester D (2024) On population-based structural health monitoring for bridges: Comparing similarity metrics and dynamic responses between sets of bridges. *Mechanical Systems and Signal Processing* 216, 111501.
- Cheyne E, Jakobsen JB and Snæbjörnsson J (2017) Damping estimation of large wind-sensitive structures. *Procedia Engineering* 199, 2047–2053.
- Cicirello A (2024) Physics-enhanced machine learning: A position paper for dynamical systems investigations. *arXiv preprint arXiv:2405.05987*.
- Dupuy D, Odier N and Lapeyre C (2024) Using graph neural networks for wall modeling in compressible anisothermal flows. *Data-Centric Engineering* 5, e10.
- Dupuy D, Odier N, Lapeyre C and Papadogiannis D (2023) Modeling the wall shear stress in large-eddy simulation using graph neural networks. *Data-Centric Engineering* 4, e7.
- Duthé G, Abdallah I, Barber S and Chatzi E (2023) Graph neural networks for aerodynamic flow reconstruction from sparse sensing. *arXiv preprint arXiv:2301.03228*.
- Duthé G, de Nolasco Santos F, Abdallah I, Réthore P-É, Weijtjens W, Chatzi E and Devriendt C (2023) Local flow and loads estimation on wake-affected wind turbines using graph neural networks and pywake. *Journal of Physics: Conference Series* 2505 (1), 012014.
- Farrar CR and Worden K (2012) *Structural Health Monitoring: A Machine Learning Perspective*. John Wiley & Sons.
- Fukunaga K (2013) *Introduction to Statistical Pattern Recognition*. Elsevier.
- Gallet A, Rigby S, Tallman T, Kong X, Hajirasouliha I, Liew A, Liu D, Chen L, Hauptmann A and Smyl D (2022) Structural engineering from an inverse problems perspective. *Proceedings of the Royal Society A* 478 (2257), 20210526.
- Gilmer J, Schoenholz SS, Riley PF, Vinyals O and Dahl GE (2017) Neural message passing for quantum chemistry. *International conference on machine learning*, 1263–1272.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning* [<http://www.deeplearningbook.org>]. MIT Press.
- Gosliga J, Gardner P, Bull L, Dervilis N and Worden K (2021) Foundations of population-based shm, part ii: Heterogeneous populations—graphs, networks, and communities. *Mechanical Systems and Signal Processing* 148, 107144.
- Gosliga J, Hester D, Worden K and Bunce A (2022) On population-based structural health monitoring for bridges. *Mechanical Systems and Signal Processing* 173, 108919.
- Gres S, Mandler A, Jacobsen N-J, Andersen P and Döhler M (2022) Statistical damage detection and localization with mahalnobis distance applied to modal parameters. *IOMAC 2022-9th International Operational Modal Analysis Conference*, 1–8.
- Greš S, Tatsis KE, Dertimanis V and Chatzi E (2023) Low-rank approximation of hankel matrices in denoising applications for statistical damage diagnosis of wind turbine blades. *Mechanical Systems and Signal Processing* 197, 110391.
- Hamilton W, Ying Z and Leskovec J (2017) Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems* 30.

- Hasan MDA, Ahmad ZAB, Leong MS and Hee LM** (2018) Enhanced frequency domain decomposition algorithm: A review of a recent development for unbiased damping ratio estimates. *Journal Of Vibroengineering* 20 (5), 1919–1936. <https://doi.org/10.21595/jve.2018.19058>.
- Haywood-Alexander M, Liu W, Bacsa K, Lai Z and Chatzi E** (2023) Discussing the spectra of physics-enhanced machine learning via a survey on structural mechanics applications. *arXiv preprint arXiv:2310.20425*.
- Hou R and Xia Y** (2021) Review on the new development of vibration-based damage identification for civil engineering structures: 2010–2019. *Journal of Sound and Vibration* 491, 115741.
- Jian X** (2025 August) *Jxdengineer/modalgnn_freq_domain: Paper_submission* (version 1). *Zenodo*. <https://doi.org/10.5281/zenodo.16779978>.
- Jian X, Lai Z, Bacsa K, Fu Y, Koh CG, Sun L, Wieser A and Chatzi E** (2024) A robotic automated solution for operational modal analysis of bridges with high-resolution mode shape recovery. *Journal of Structural Engineering* 150 (8), 04024081.
- Jian X, Xia Y and Sun L** (2022) Indirect identification of bridge frequencies using a four-wheel vehicle: Theory and three-dimensional simulation. *Mechanical Systems and Signal Processing* 177, 109155.
- Kamariotis A, Chatzi E and Straub D** (2022) Value of information from vibration-based structural health monitoring extracted via bayesian model updating. *Mechanical Systems and Signal Processing* 166, 108465.
- Karakostas C, Quaranta G, Chatzi E, Zülfiqar AC, Çetindemir O, De Roeck G, Döhler M, Limongelli MP, Lombaert G, Apaydin NM, et al.** (2024) Seismic assessment of bridges through structural health monitoring: A state-of-the-art review. *Bulletin of Earthquake Engineering* 22 (3), 1309–1357.
- Kipf TN and Welling M** (2016) Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Lai Z, Liu W, Jian X, Bacsa K, Sun L and Chatzi E** (2022) Neural modal ordinary differential equations: Integrating physics-based modeling with neural ordinary differential equations for modeling high-dimensional monitored structures. *Data-Centric Engineering* 3, e34.
- Lei X, Siringoringo DM, Dong Y and Sun Z** (2023) Interpretable machine learning methods for clarification of load-displacement effects on cable-stayed bridge. *Measurement* 220, 113390.
- Limongelli MP, Gentile C, Biondini F, di Prisco M, Ballio F, Zonno G, Borlenghi P, Bianchi S, Capacci L, Anghileri M, et al.** (2024) Bridge structural monitoring: The Lombardia regional guidelines. *Structure and Infrastructure Engineering* 20 (4), 461–484.
- Moreu F, Li X, Li S and Zhang D** (2018) Technical specifications of structural health monitoring for highway bridges: New chinese structural health monitoring code. *Frontiers in Built Environment* 4. <https://doi.org/10.3389/fbuil.2018.00010>.
- Mylonas C** (2021) *Machine Learning for Structural Health Assessment under Uncertainty. With Applications in Wind Energy [Doctoral Thesis]*. ETH Zurich. ETH Zurich. <https://doi.org/10.3929/ethz-b-000511551>
- Persson P-O and Strang G** (2004) A simple mesh generator in matlab. *SIAM Review* 46 (2), 329–345.
- Reynders E** (2012) System identification methods for (operational) modal analysis: Review and comparison. *Archives of Computational Methods in Engineering* 19 (1), 51–124. <https://doi.org/10.1007/s11831-012-9069-x>.
- Rossi E, Kenlay H, Gorinova MI, Chamberlain BP, Dong X and Bronstein MM** (2022) On the unreasonable effectiveness of feature propagation in learning on graphs with missing node features. *Learning on Graphs Conference*, 1–11.
- Sanchez-Lengeling, B., Reif, E., Pearce, A., & Wiltchko, A. B.** (2021). *A Gentle Introduction to Graph Neural Networks* [<https://distill.pub/2021/gnn-intro>]. *Distill*. <https://doi.org/10.23915/distill.00033>
- Stoura CD, Dertimanis VK, Hoelzl C, Kossmann C, Cigada A, Chatzi EN, et al.** (2023) A model-based bayesian inference approach for on-board monitoring of rail roughness profiles: Application on field measurement data of the swiss federal railways network. *Structural Control and Health Monitoring* 2023.
- Sun L, Shang Z, Xia Y, Bhowmick S and Nagarajaiah S** (2020) Review of bridge structural health monitoring aided by big data and artificial intelligence: From condition assessment to damage detection. *Journal of Structural Engineering* 146 (5), 04020073.
- Tsialiamanis G, Dervilis N, Wagg DJ and Worden K** (2023) Towards a population-informed approach to the definition of data-driven models for structural dynamics. *Mechanical Systems and Signal Processing* 200, 110581.
- Tsialiamanis G, Mylonas C, Chatzi EN, Wagg DJ, Dervilis N and Worden K** (2022) On an application of graph neural networks in population-based shm. *Data Science in Engineering, Volume 9: Proceedings of the 39th IMAC, A Conference and Exposition on Structural Dynamics* 2021, 47–63.
- Tsialiamanis G, Sbarufatti C, Dervilis N and Worden K** (2024) On a meta-learning population-based approach to damage prognosis. *Mechanical Systems and Signal Processing* 209, 111119.
- Van Der Maaten L, Postma EO, van den Herik HJ, et al.** (2009) Dimensionality reduction: A comparative review. *Journal of Machine Learning Research* 10 (13), 66–71.
- Veličković P, Cucurull G, Casanova A, Romero A, Lio P and Bengio Y** (2017) Graph attention networks. *arXiv preprint arXiv:1710.10903*.
- Wu Z, Pan S, Chen F, Long G, Zhang C and Philip SY** (2020) A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems* 32 (1), 4–24.
- Zhou J, Cui G, Hu S, Zhang Z, Yang C, Liu Z, Wang L, Li C and Sun M** (2020) Graph neural networks: A review of methods and applications. *AI Open* 1, 57–81.

Cite this article: Jian X, Xia Y, Duthé G, Bacsa K, Liu W and Chatzi E (2025). Using graph neural networks and frequency domain data for automated operational modal analysis of populations of structures. *Data-Centric Engineering*, 6, e45. doi:10.1017/dce.2025.10023