

Temperature-Robust Structural Damage Detection via Unsupervised Domain Adaptation

Soyeon Park ^a, Junho Song ^{a*}

^a Department of Civil and Environmental Engineering, Seoul National University, 1 Gwanak-ro, Gwanak-gu, Seoul, Republic of Korea

*Corresponding author: junhosong@snu.ac.kr

***Keywords :** structural health monitoring, temperature, domain adaptation, unsupervised learning, CVAE

1. Introduction

Structural Health Monitoring (SHM) is essential for detecting damage in nuclear power plant structures, but its credibility is often compromised by environmental variability. Temperature changes alter material properties and modal responses, which can mask or mimic structural damage and lead to false alarms and missed detections. Conventional compensation methods such as regression, cointegration, or Principal Component Analysis (PCA) assume linear and stationary relations, making them inadequate under multi-sensor conditions.

To address this issue, we propose an unsupervised learning framework for temperature-robust damage detection. The method combines frequency patching to capture localized spectral shifts, a Conditional Variational Autoencoder (CVAE) to learn compact latent features of healthy states, and Maximum Mean Discrepancy (MMD)-based domain adaptation to align distributions across different temperatures. A temperature-aware scoring scheme is then applied to provide consistent anomaly indices. This approach enables reliable detection of structural damage without requiring labeled damage data.

2. Methodology

The proposed method consists of four major stages: frequency-domain preprocessing, generative modeling with CVAE, latent alignment with MMD, and temperature-aware scoring.

2.1 Frequency patching

To extract damage-sensitive but temperature-robust features, the raw time-domain responses must first be converted into a spectral representation. Given a multi-channel time-domain response window $x \in \mathbb{R}^{L \times C}$, where L is the number of samples, and C is the number of sensor channels, a Hann window $w \in \mathbb{R}^L$ is applied. The real Fourier Transform is then computed as

$$X(f, c) = rFFT[w \odot x(:, c)], f = 0, 1, \dots, F_+, \quad (1)$$

where $\odot$ denotes element-wise multiplication, $X(f, c)$ is the complex coefficient at frequency index f for channel c and $F_+ = [L/2]$ is the number of positive frequency bins.

The log-magnitude spectrum is then obtained as

$$M(f, c) = \log(1 + |X(f, c)|), \quad (2)$$

The frequency axis is divided into non-overlapping bands of width B . Each patch is defined as

$$X[p, :, :] = M[pB:(p+1)B, :], p = 0, \dots, P-1, \quad (3)$$

with $P = [F_+/B]$. The resulting tensor $X \in \mathbb{R}^{P \times B \times C}$ captures the localized frequency behavior.

2.2 Conditional Variational Autoencoder (CVAE)

To build a compact generative model representing responses of a healthy structure, we employ a CVAE. The encoder defines an approximate posterior

$$q_\theta(z|X) = \mathcal{N}\left(z; \mu_\theta(X), \text{diag}(\sigma_\theta^2(X))\right), \quad (4)$$

where $z \in \mathbb{R}^d$ is the latent vector of dimension d , and $\mu_\theta(X) \in \mathbb{R}^d$ and $\sigma_\theta^2(X) \in \mathbb{R}^d$ are the encoder outputs parameterized by weights θ . The decoder reconstructs patches as

$$p_\phi(X|z) = \mathcal{N}(\hat{X}; f_\phi(z), \sigma^2 I), \quad (5)$$

with parameter ϕ .

The training objective is the evidence lower bound (ELBO):

$$\mathcal{L}_{ELBO}(X) = \|X - \hat{X}\|_1 + \beta \cdot KL[q_\theta(z|X) \parallel \mathcal{N}(0, I)], \quad (6)$$

where β is a weighting factor. The KL divergence term has closed form:

$$KL[q_\theta(z|X) \parallel \mathcal{N}(0, I)] = \frac{1}{2} \sum_{k=1}^d (\mu_k^2 + \sigma_k^2 - \log \sigma_k^2 - 1). \quad (7)$$

where μ_k and σ_k^2 denote the k -th components of $\mu_\theta(X)$ and $\sigma_\theta^2(X)$.

2.3 Domain adaptation with MMD

Since data of healthy structural responses under different temperature ranges belong to different distributions, direct application of the CVAE may still entangle temperature variation with structural characteristics. To overcome this, we introduce MMD for unsupervised domain adaptation. Let $Z_S = \{z_i^S\}_{i=1}^b$ and $Z_T = \{z_j^T\}_{j=1}^b$ denote mini-batches of latent vectors from source (C2: wide range temperature, healthy) and target (C1: narrow range temperature, healthy) domains. The MMD is then defined as

$$MMD^2(Z_S, Z_T) = \frac{1}{b(b-1)} \sum_{i \neq i'} k(z_i^S, z_{i'}^S) + \frac{1}{b(b-1)} \sum_{j \neq j'} k(z_j^T, z_{j'}^T) - \frac{2}{b^2} \sum_{ij} k(z_i^S, z_j^T), \quad (8)$$

where the kernel is

$$k(u, v) = \sum_{m=1}^M \exp\left(-\frac{\|u - v\|_2^2}{2\sigma_m^2}\right). \quad (9)$$

The Stage-B training objective combines reconstruction, KL regularization, and MMD alignment: $\mathcal{L} = \mathbb{E}_{X \in (C1 \cup C2)} [\mathcal{L}_{ELBO}(X)] + \lambda_{MMD} \cdot MMD^2(Z_S, Z_T)$ (10) where λ_{MMD} is the adaption weight.

2.4 Temperature-aware scoring

Finally, to detect anomalies in a temperature-consistent manner, reconstruction residuals are normalized by temperature-specific statistics. For a test sample, the residual is

$$r(x) = \|X - \hat{X}\|_1. \quad (11)$$

Residuals of the healthy structural responses are collected per temperature bin $b(T)$, with mean $\mu_{b(T)}$ and standard deviation $\sigma_{b(T)}$. The anomaly score is defined as

$$S(x, T) = \frac{r(x) - \mu_{b(T)}}{\sigma_{b(T)}}. \quad (12)$$

3. Numerical simulations

3.1 Data

A numerical analysis was conducted on a 9-story shear building idealization of the benchmark example by Ohtori et al. (2004). The model was subjected to a small-amplitude, broadband, zero-mean stochastic base acceleration used as a proxy for ambient vibration. During the implementation, the excitation is modeled as a discrete-time white Gaussian base-acceleration sequence with its variance calibrated to keep responses within the linear-elastic range, for consistency with operational modal-analysis practice. The reference benchmark represents a steel frame with 198 degrees of freedom (DOFs), consisting of 21 nodes and 28 structural elements (columns and beams). The structure was simplified into a 9-DOF shear building by calculating equivalent mass and stiffness matrices for each floor, as illustrated in Figure 1.

Structural damage was introduced by reducing the stiffness of selected column members, which translates into a reduction in story stiffness in the 9-DOF model.




Fig. 1. Schematics of a 9-story shear building

Temperature effects were modeled by adjusting the stiffness of each story according to the empirical relationship

$$f_T \approx f_0 \cdot \sqrt{1 - \beta_E \cdot \Delta T}, \quad (13)$$

where f_T is the natural frequency at temperature T , f_0 is the baseline temperature, β_E denotes the temperature-elastic modulus coefficient, and ΔT is the deviation from the baseline temperature.

By combining these two mechanisms (i) stiffness reduction for damage and (ii) temperature-dependent frequency shifts, three datasets were generated:

- C1 (Target domain): Healthy responses under a 20°C
- C2 (Source domain): Healthy responses under a wide temperature range (−10~40°C)
- C3 (Damaged): Responses with imposed stiffness reductions at various floors under multiple temperatures

Figure 2. shows that the fundamental frequency is influenced by both temperature and damage degree, decreasing systematically as temperature rises and stiffness reduction increases.




Fig. 2. Contour plot of the fundamental frequency under various temperature and damage degree

3.2 Results

To quantify temperature-robust damage detection, we compared the results by a conventional PCA-only baseline with those by the proposed method.




Fig. 3. Score distributions: (a) PCA-only, and (b) proposed temperature-invariant method

Figure 3 provides normalized histograms where the x-axis is the temperature-aware z -score of the reconstruction residual, and the y-axis is the probability density; the vertical dashed line marks the 95th percentile of healthy scores. The results by “PCA-only” in Figure 3(a) show substantial overlap of healthy (C1+C2) and damaged (C3) near the threshold, indicating residual temperature leakage and weaker separability. In contrast, the results by the proposed method in Figure 3(b) displays a clear left–right separation: healthy samples concentrate left of the threshold, while damaged shift right, yielding higher detection performance and more stable behavior across temperatures.

4. Conclusions

We introduced a temperature-robust damage detector that turns responses into frequency patches, learns a healthy manifold with an CVAE, and uses MMD to make the latent space temperature-invariant; decisions rely on a temperature-aware z -score. Compared with a PCA-only baseline, it shows clearer healthy–damage separation and a flatter healthy false positive ratio ($\approx 5\%$) without using damaged data for training. Further research is underway to validate using seismic responses and real-bridge data, and embed physics (mode-band priors, temperature–stiffness relations, damping/peak constraints) to derive a damage index that is sensitive to damage severity.

REFERENCES

- [1] Maghareh, A., Barnes, M. J., and Sun, Z. *Evaluation of the 9-story benchmark building shear model*. University of Notre Dame, 2012.
- [2] Ohtori, Y., Christenson, R. E., Spencer, B. F., and Dyke, S. J. "Benchmark control problems for seismically excited nonlinear buildings." *ASCE Journal of Engineering Mechanics* 130, no. 4 (2004): 366–385.
- [3] Pol, A. A., Berger, V., Germain, C., Cerminara, G., and Pierini, M. "Anomaly detection with conditional variational autoencoders." In *2019 18th IEEE International*

Conference on Machine Learning and Applications (ICMLA), 1651–1657. IEEE, 2019.

- [4] Luo, J., Huang, M., and Lei, Y. "Temperature effect on vibration properties and vibration-based damage identification of bridge structures: A literature review." *Buildings* 12, no. 8 (2022): 1209.