



ARTICLE

Improving Stage-Transition Detection in Long-Term Guided-Wave Monitoring through Autoencoder Regularization

Boyang Zhang¹, Kang Gao², Jianan Gu³, Li Ai⁴ and Yang Geng^{3,*}

¹Department of Electrical and Computer Engineering, Boston University, Boston, MA, USA

²Department of Electrical and Computer Engineering, Michigan State University, East Lansing, MI, USA

³Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL, USA

⁴Department of Electrical and Computer Engineering, University of Texas Rio Grande Valley, Edinburg, TX, USA

*Corresponding Author: Yang Geng. Email: yang.geng@ufl.edu

Received: 01 April 2026; Accepted: 02 June 2026; Published: 23 October 2026

ABSTRACT: Long-term guided-wave structural health monitoring accumulates large volumes of sensing data, but dense damage labels are rarely available in field deployment. Unsupervised reconstruction-based detection methods, which identify damage through reconstruction discrepancy, are therefore appealing for their ability to learn directly from unlabeled data. However, environmental and operational variations and mixed adjacent-stage samples in transition windows both degrade unsupervised detection performance. Regularization offers a principled way to improve the robustness of unsupervised detection against such confounding variations, but its effect on fine-grained, transition-level detection has not been systematically studied. This paper presents a controlled empirical study to evaluate various regularization methods for unsupervised transition-level damage detection under uncontrolled environmental conditions. Using the public Scientific Data benchmark, four regularization configurations, namely unregularized, dropout, weight decay, and sparsity, are compared under an identical architecture and training protocol. Detection performance is evaluated at adjacent stage transitions, which reflect the practical question of whether a new structural state differs from the immediately preceding one. Pearson correlation coefficient (PCC) trajectories and area under the receiver operating characteristic curve (AUROC) serve as the primary metrics, and a transition is counted as detected only when both $\text{AUROC} \geq 0.85$ and average PCC gap between two states ≥ 0.010 are satisfied. Under this criterion, the unregularized autoencoder detects 2 of 12 adjacent transitions, while dropout, weight decay, and sparsity detect 5, 6, and 7 transitions, respectively, with best mean AUROC values of 0.9071, 0.9249, and 0.9288. Analysis of representative difficult cases reveals two distinct failure regimes: weak structural contrast between adjacent stages and environment-dominated signal masking. These results demonstrate that regularization alone, without architectural modification, can materially increase transition-level detection coverage in long-term uncontrolled guided-wave monitoring.

KEYWORDS: Guided-wave structural health monitoring; unsupervised damage detection; autoencoder regularization; uncontrolled environment; adjacent stage-transition detection; reconstruction correlation

1 Introduction

Structural health monitoring (SHM) is essential for ensuring the safety and reliability of critical infrastructure and mechanical systems such as bridges, pipelines, turbines, and composite panels. In many of these applications, structural damage develops gradually or is embedded within complex materials and boundary conditions, making periodic human inspection alone insufficient for timely detection [1]. Data-driven approaches have therefore become increasingly important in modern SHM practice. These

approaches collect large volumes of sensing data over extended periods and detect structural damage through systematic computational analysis rather than manual assessment. Among the available sensing modalities, guided-wave monitoring is particularly suitable for this purpose because it enables long-range inspection with high sensitivity to local structural changes [2]. These properties have made guided waves one of the most widely studied modalities in SHM [3–5].

To effectively detect structural damage from guided-wave data, a variety of signal processing techniques have been developed. Conventional approaches include baseline subtraction [6], damage index computation [7], which compare current measurements against a stored reference to identify deviations indicative of structural change. To improve robustness against environmental variation, temperature compensation [8], scale-transform correction [9], dynamic time warping alignment [10], and singular-value-decomposition stabilization [11] have also been proposed. However, these methods generally depend on handcrafted processing rules and explicit reference baselines, limiting their adaptability to complex and nonstationary monitoring environments.

Learning-based methods offer a more flexible alternative by extracting damage-sensitive features directly from data. Supervised deep-learning approaches have achieved strong detection accuracy when sufficient labeled data are available [12,13]. In long-term monitoring, however, damage labels are scarce, expensive to obtain, and frequently unavailable for rare or gradually evolving fault states, making fully supervised learning impractical for most field deployments. Semi-supervised strategies relax the labeling requirement by training on data assumed to represent only the reference structural condition [14], but this assumption is difficult to guarantee in practice: uncontaminated baselines are hard to maintain in uncontrolled environments, and the definition of reference operation may drift over time. These limitations motivate unsupervised learning, which operates without damage annotations and is therefore more compatible with realistic long-term monitoring workflows.

Among unsupervised approaches, reconstruction-based methods using autoencoders have become a widely adopted framework for identifying structural-state deviations from unlabeled monitoring data [15,16]. In the conventional binary reference/deviation interpretation, an autoencoder trained on reference data learns dominant baseline patterns and reconstructs less represented deviation patterns with lower fidelity. In the present study, we adapt this reconstruction principle to adjacent damage-stage transition detection: the relevant contrast is not a binary reference-vs.-deviation decision, but the reconstruction difference between a preceding structural stage and its immediately subsequent, more severe stage.

Despite this conceptual simplicity, reconstruction-based unsupervised detection faces two practical challenges in long-term guided-wave monitoring. First, if the model capacity is excessive, the autoencoder can learn to faithfully reconstruct waveform patterns from both adjacent stages, reducing stage-transition contrast and thereby degrading detection performance. Second, environmental and operational conditions (EOCs), including temperature, humidity, precipitation, sensor drift, and installation shifts, introduce substantial variability into the reference-stage measurements [17]. This variability degrades reconstruction consistency and can produce reconstruction changes comparable to those caused by structural-stage progression, making adjacent-stage differences difficult to distinguish from environmental disturbances.

Regularization provides a principled means to address both challenges. Previous studies have demonstrated that regularization and training protocol strongly affect the performance of unsupervised autoencoders in long-term guided-wave monitoring [18–20]. By constraining model capacity or activation behavior during training, regularization intentionally limits what the autoencoder can learn, reducing the risk that it reconstructs all adjacent-stage waveform variations equally well while improving robustness to environmental noise. Common strategies include dropout, which randomly deactivates hidden units during training to prevent co-adaptation; weight decay, which penalizes large parameter values to encourage simpler

learned representations; and sparsity constraints, which enforce compact latent activations to concentrate the model on the most informative signal components.

Although the effectiveness of regularization has been recognized in individual studies, a comprehensive empirical investigation that systematically compares different regularization strategies under one fixed pipeline remains lacking. In particular, no existing work evaluates regularization in terms of adjacent stage-transition detection, where the goal is to determine whether a newly observed structural state can be distinguished from the immediately previous stage. This type of detection directly corresponds to the gradual damage progression encountered in practical maintenance scenarios and is more demanding than the commonly used pooled comparison between extreme healthy and severe-damage endpoints. Moreover, the effect of regularization under uncontrolled environmental and operational conditions, where EOC effects are realistic and unmitigated, remains insufficiently studied despite its direct relevance to field deployment. Importantly, the problem studied here is not conventional healthy-vs.-damaged detection, but reconstruction-based adjacent damage-stage transition detection. The goal is to determine whether a newly emerged structural state can be distinguished from its immediately preceding state through reconstruction behavior under a local, unlabeled transition window.

To address these gaps, this paper presents a controlled empirical study of regularization for unsupervised damage detection in long-term guided-wave monitoring under uncontrolled environments. This work does not aim to introduce a new autoencoder architecture or a new regularization formula. Instead, it addresses a deployment-oriented methodological question that remains underexplored in long-term guided-wave SHM: how standard regularization choices affect unsupervised adjacent damage-stage transition detection under uncontrolled environmental and operational variability. The reconstruction pipeline is held fixed and only the regularization mechanism is varied, so that performance differences can be attributed to regularization rather than architecture changes. Three regularization schemes, dropout, weight decay, and sparsity, are compared against a vanilla autoencoder benchmark under identical preprocessing and training settings. Using the Pearson correlation coefficient (PCC) and area under the receiver operating characteristic curve (AUROC) as detection indicators, under a fixed standard, we evaluate each regularization scheme by its ability to detect adjacent damage-stage transitions, with particular attention to whether subtle progression between early and mild damage states can be captured. In addition to successful detection examples, difficult cases are explicitly analyzed to reflect deployment realism.

The contributions of this paper are summarized as follows:

1. A controlled comparison of dropout, weight decay, and sparsity is conducted under a fixed autoencoder pipeline, allowing the effect of regularization to be examined without introducing confounding architectural changes.
2. An adjacent damage-stage transition detection protocol is used to evaluate whether a newly emerged, more severe structural state can be distinguished from its immediately preceding state using local transition windows.
3. Practical detection behavior is analyzed through pair coverage, threshold sensitivity, hyperparameter sensitivity, path-wise behavior, and representative failure regimes, providing deployment-oriented guidance for regularized reconstruction-based monitoring under uncontrolled environments.

2 Related Work

2.1 Methods for Guided-Wave Damage Detection under Environmental Variability

A central challenge in guided-wave structural health monitoring is that environmental and operational variations can alter signal characteristics as significantly as actual structural damage [1]. In long-term field

deployment, temperature fluctuation, humidity change, precipitation, sensor drift, and installation shift all contribute to waveform distortion, making it difficult to attribute observed signal changes to structural causes alone [17]. Recent Bayesian studies further show that guided-wave interpretation can be affected by sensor attachment uncertainty and by temperature-dependent uncertainty in Lamb-wave model calibration [21,22]. Two broad categories of methods have been developed to address this challenge: signal-level environmental compensation and data-driven unsupervised detection.

Signal-level compensation methods aim to remove or reduce environmental effects before damage assessment. Early work established temperature compensation strategies for guided-wave SHM [8]. Scale-transform correction subsequently provided a principled approach to ultrasonic temperature compensation [9], while dynamic time warping was applied to align thermally distorted signals [10]. Robust singular-value-decomposition stabilization further improved detection under complex environmental conditions [11]. These methods improve signal stability under controlled or slowly varying conditions, but they rely on explicit compensation logic tailored to specific disturbance types and do not directly address how an unsupervised model should be configured for long-term monitoring under broad and uncontrolled variability.

Data-driven methods offer a complementary approach by learning signal structure directly from data without requiring explicit disturbance models. Early data-driven work in guided-wave SHM adopted principal component analysis (PCA) and its variants for unsupervised feature extraction and anomaly detection. As modeling demands grew, deep learning methods further expanded the capacity for nonlinear representation learning in SHM. Deep-learning-enhanced PCA bridged classical low-dimensional modeling with learned representations for structural damage assessment [23]. GAN-based anomaly modeling demonstrated the value of generative representation learning in label-scarce monitoring settings [16]. Robust state-space learning showed that unsupervised detection under mixed or contaminated data requires models with explicit robustness mechanisms [24]. Related guided-wave studies investigated sparse representations under varying environmental and operational conditions [25], unsupervised feature learning for delamination-sensitive representations [26], unsupervised damage localization [27], and detection performance under environmental and temperature variation [17]. Among these methods, autoencoder-based reconstruction has become a widely adopted unsupervised approach for damage-sensitive monitoring, owing to its simple training procedure and direct reconstruction-quality scoring [12,15]. It has been shown that stacked autoencoders can extract damage-sensitive features effectively from monitored structures [28]. Convolutional autoencoders have also been applied to detect damage directly from raw vibration signals without manual feature engineering [29]. Variational autoencoders further enable baseline-free structural damage identification by learning latent representations of structural responses [30].

Building on those autoencoder-based detection methods, recent studies have shown that regularization strategy, training reference selection, and training protocol strongly affect unsupervised detection performance and can be more influential than architectural modifications [18,19]. Weight-decay-oriented autoencoder benchmarking further confirmed that explicit regularization design is necessary for stable guided-wave monitoring and provided comparative evidence for one specific regularization type [20]. However, these studies each focused on a particular aspect of autoencoder configuration rather than systematically comparing multiple regularization mechanisms under one fixed pipeline.

Motivated by these findings, this paper presents a systematic empirical study of regularization for unsupervised autoencoder-based damage detection in long-term guided-wave monitoring. Unlike previous studies that focused on a particular aspect of autoencoder configuration [18–20], this work conducts a controlled comparison across multiple regularization mechanisms under one fixed pipeline, uses adjacent

stage-transition detection as the primary evaluation target, and provides explicit interpretation of representative failure modes. The contribution is therefore a deployment-oriented evaluation and interpretation protocol for regularization choices in long-term uncontrolled guided-wave monitoring.

2.2 Benchmark Datasets for Guided-Wave SHM

The availability of public benchmark datasets has been essential for advancing reproducible evaluation in guided-wave SHM. Open Guided Waves established a public platform for reproducible ultrasonic guided-wave comparison [2]. Subsequent datasets extended benchmark coverage to more complex composite configurations [3] and to full wavefield measurements across bonded and debonded states [4]. Most relevant to the present study, the long-term guided-wave monitoring dataset built by Kang et al. [5] captures guided-wave signals under uncontrolled and dynamic environmental conditions. This dataset covers an aluminum plate monitored by a pitch-catch piezoelectric transducer network over a period spanning from April 2021 to October 2022. The structural condition progresses through 13 stages: one healthy baseline and 12 levels of increasing damage, ranging from a small attached mass to progressively larger dents and through-holes with accompanying dent damage. Each stage transition represents a gradual and controlled change in structural condition, and the monitoring duration within each stage spans weeks to months under natural outdoor weather. The dataset also provides timestamped environmental metadata (temperature, humidity, brightness, air pressure, and weather tags) alongside guided-wave measurements and damage-stage labels. These characteristics make this dataset particularly suitable for the present empirical study: the 13-stage progression provides sufficient granularity to evaluate adjacent-transition detection rather than only extreme-state separation, the long monitoring duration exposes the full range of seasonal and weather-driven variability, and the gradual damage evolution closely resembles the incremental deterioration encountered in real structural aging. Because it captures both the complexity of outdoor environmental conditions and a fine-grained structural progression, this dataset serves as the experimental foundation for the present study.

3 Methods

In this work, we investigate how regularization improves unsupervised adjacent stage-transition detection for long-term guided-wave monitoring under uncontrolled conditions, a task that directly reflects practical maintenance decisions where structural damage progresses gradually. In this section, we detail the dataset and signal preparation (Section 3.1), the reconstruction-based detection framework (Section 3.2), the regularization strategies (Section 3.3), and the evaluation protocol (Section 3.4).

The task studied here is ordered adjacent damage-stage transition detection rather than conventional binary healthy-vs.-damaged detection. For each adjacent transition, we construct a five-day transition window consisting of 4.5 days of measurements from the preceding structural state and 0.5 days of measurements from the subsequent, more severe structural state, as summarized in Table 1. The autoencoder is trained in an unsupervised reconstruction manner using the unlabeled samples in this transition window. Damage-stage labels are not used in the reconstruction loss, model optimization, classifier learning, or threshold tuning. They are used only after training to compute transition-level evaluation metrics, including AUROC, the average PCC gap between the preceding and subsequent stages, and detected-pair coverage.

This setting differs from conventional binary healthy-vs.-damaged detection. The objective is not merely to determine whether damage exists, but to evaluate whether a newly emerged, more severe damage state can be distinguished from the immediately preceding state through reconstruction behavior. This problem definition is central to the present study: the subsequent-stage samples are not treated as globally abnormal samples, but as locally underrepresented, newly emerged structural responses relative to the immediately preceding stage. Because the transition window contains substantially more samples from the preceding

state than from the subsequent state, the empirical training distribution is dominated by the preceding structural condition. Regularization is therefore expected to preserve reconstruction contrast by preventing the autoencoder from fitting all stage-specific waveform variations equally well.

Table 1: Transition-window protocol used for adjacent stage-transition detection. Approximate sample counts are based on the released public dataset, which keeps one measurement approximately every 176 s after 20-fold downsampling. Each retained measurement contains one signal for path 6-1, corresponding to approximately 491 single-path samples per day.

Window Component	Duration	Approx. Path-6-1 Samples	Role
Preceding/reference stage	4.5 days	~2209	Dominant/reference state
Subsequent/newly emerged stage	0.5 day	~245	More severe newly emerged state
Total transition window	5.0 days	~2454	Unsupervised autoencoder training/evaluation window

3.1 Dataset and Signal Preparation

The dataset used in this study is the public long-term guided-wave monitoring dataset published by Kang et al. [5], which was collected under uncontrolled and dynamic environmental conditions. The monitored structure is an aluminum plate (53 cm × 53 cm × 3 mm) equipped with a pitch-catch piezoelectric transducer (PZT) network. The PZT network consists of six sensors: sensors 1 through 4 serve as receivers, while sensors 5 and 6 serve as transmitters, yielding eight guided-wave paths: 5-1, 5-2, 5-3, 5-4, 6-1, 6-2, 6-3, and 6-4. All main experiments are conducted on path 6-1 as a relatively stable single-path benchmark. The dataset validation shows that long-term sensor drift is path-dependent and that path 6-1 exhibits milder drift than several other representative paths [5]. We therefore use path 6-1 as a conservative single-channel setting to reduce confounding from severe path-dependent sensor degradation and to more clearly isolate the effect of regularization on adjacent stage-transition detection. This choice does not imply that path 6-1 is universally optimal; rather, it defines the controlled single-path setting used before expanded path-level and multi-channel analyses. The path-selection rationale is summarized in Fig. 1.

The transmitted signal is a 1 ms linear chirp excitation spanning 5 to 350 kHz, sampled at 1 MHz, with measurements recorded every 8.6 s. The public benchmark release applies 20-fold temporal downsampling (approximately one retained measurement every 176 s) and trims each guided-wave signal to the first 2000 samples. We use this released format directly. Monthly data files (`measurements_YYYY_MM.pickle`) store a guided-wave tensor of dimension $N \times 8 \times 2000$, where N is the number of measurements in that month. Each record also includes environmental metadata: timestamp, temperature, humidity, brightness, air pressure, damage tag, and weather tag.

In the released dataset, damage tag 0 corresponds to the healthy condition and tags 1 through 12 correspond to progressively damaged conditions [5]. For adjacent-transition analysis, we re-index these into a 13-stage sequence $D_1, D_2, \dots, D_{13}$, where D_1 corresponds to damage tag 0 (the healthiest baseline) and D_k corresponds to damage tag $k - 1$ for $k = 2, \dots, 13$. Weather tag 0 denotes fair weather, while tags 1 through 5 correspond to precipitation-related categories. The stage labels are used only for post-training evaluation

of adjacent transitions and are not provided to the autoencoder during reconstruction training. All model training, evaluation, and representative plots in this paper are based on path 6-1.




Figure 1: Schematic rationale for using path 6-1 as the main single-path benchmark. The public dataset contains eight pitch-catch paths from transmitters 5/6 to receivers 1–4. Because the dataset validation reports path-dependent sensor drift, the main experiments use the relatively stable path 6-1 to reduce severe path-dependent confounding and isolate the effect of regularization.

3.2 Autoencoder-Based Reconstruction Detection

We adopt an autoencoder (AE) as the reconstruction backbone for unsupervised adjacent-stage transition detection. The approach relies on a reconstruction consistency assumption within each local transition window: because the window contains many more measurements from the preceding stage than from the subsequent stage, the autoencoder is expected to reconstruct dominant preceding-stage waveform patterns more faithfully than newly emerged subsequent-stage components. This assumption is naturally suited to adjacent-transition detection, where the objective is to preserve reconstruction contrast between two ordered structural stages rather than to assign binary condition labels.

Let $x_i \in \mathbb{R}^M$ denote a guided-wave signal with M samples. The autoencoder maps each input through an encoder f_θ and a decoder g_ϕ :

$$z_i = f_\theta(x_i), \quad \hat{x}_i = g_\phi(z_i), \quad (1)$$

where z_i is the latent representation and $\hat{x}_i$ is the reconstructed signal. The training objective minimizes the mean squared reconstruction error:

$$\mathcal{L}_{\text{rec}} = \frac{1}{N} \sum_{i=1}^N \|x_i - \hat{x}_i\|_2^2. \quad (2)$$

The model is trained without access to stage labels and produces an unsupervised reconstruction-discrepancy score for each sample based on reconstruction quality.

To quantify reconstruction quality, we use the Pearson correlation coefficient (PCC) between the original and reconstructed signals:

$$\text{PCC}_i = \frac{(x_i - \bar{x}_i)^\top (\hat{x}_i - \bar{\hat{x}}_i)}{\|x_i - \bar{x}_i\|_2 \|\hat{x}_i - \bar{\hat{x}}_i\|_2}. \quad (3)$$

A lower PCC indicates poorer reconstruction fidelity and stronger deviation from the learned dominant transition-window patterns, so the reconstruction-discrepancy score is defined as $s_i = 1 - \text{PCC}_i$. To assess whether two adjacent stages are separable, we compute the stage-wise average PCC values and their gap:

$$\bar{r}_k = \frac{1}{N_k} \sum_{i \in D_k} \text{PCC}_i, \quad \bar{r}_{k+1} = \frac{1}{N_{k+1}} \sum_{i \in D_{k+1}} \text{PCC}_i, \quad (4)$$

$$\Delta_{\text{PCC}}^{(k)} = |\bar{r}_k - \bar{r}_{k+1}|. \quad (5)$$

The PCC gap captures whether the average reconstruction behavior shifts meaningfully between adjacent stages. In addition, sample-level separability is quantified by the area under the receiver operating characteristic curve (AUROC), which measures how well the reconstruction-discrepancy scores separate the two stages.

A transition pair $D_k \rightarrow D_{k+1}$ is classified as *detected* only when both conditions are satisfied under the same model output:

$$\text{AUROC}^{(k)} \geq 0.85, \quad \Delta_{\text{PCC}}^{(k)} \geq 0.010. \quad (6)$$

The AUROC threshold ensures sufficient sample-level discrimination, while the PCC-gap threshold ensures a meaningful shift in stage-level reconstruction quality. No pair-specific threshold tuning is applied; this criterion is uniform across all transitions and all models. Full transition-wise AUROC and PCC-gap results are provided in Supplementary Table S1, confirming that the detected-pair coverage is derived from continuous pair-level metrics rather than from a binary-only summary.

In this work, the autoencoder uses a fixed architecture and training configuration across all experiments to ensure a fair comparison. The network and optimization details are summarized in [Table 2](#). For visual consistency, all representative PCC trajectory plots in [Figs. 2–4](#) are drawn from epoch-60 model outputs.

Table 2: Autoencoder network architecture and training configuration.

Parameter	Setting
Input dimension	$M = 2000$
Encoder hidden layers	$2000 \rightarrow 1024 \rightarrow 256 \rightarrow 64$
Decoder hidden layers	$64 \rightarrow 256 \rightarrow 1024 \rightarrow 2000$
Latent dimension	64
Batch size	256
Learning rate	5×10^{-4}
Maximum epochs	71
Evaluation epoch	60




Figure 2: Representative moderate-contrast transitions: the vanilla AE is already weakly informative, while regularized variants produce more evident and more stable separation. Left, $D_4 \rightarrow D_5$; right, $D_7 \rightarrow D_8$.




Figure 3: Representative baseline-ineffective transitions recovered by regularization: the vanilla AE alone does not provide sufficient discrimination, whereas regularized variants achieve clear separation. Left, $D_6 \rightarrow D_7$; right, $D_9 \rightarrow D_{10}$.




Figure 4: Representative difficult transitions. Left, weak-contrast case $D_2 \rightarrow D_3$; right, environment-dominated case $D_{10} \rightarrow D_{11}$. These cases are important for deployment realism and failure analysis.

3.3 Regularization Strategies

In reconstruction-based adjacent-transition detection, the core mechanism is that samples departing from the learned dominant waveform patterns exhibit poorer reconstruction quality. To produce useful adjacent-transition contrast, the autoencoder must reconstruct signals from the preceding stage more faithfully than signals from the subsequent stage. However, an unrestricted autoencoder with sufficient capacity can learn to reconstruct both preceding-stage patterns and newly emerged or environmentally distorted subsequent-stage patterns with comparable fidelity [18]. When the model reconstructs both stages equally well, reconstruction discrepancy becomes uniformly low and stage-transition contrast vanishes.

Regularization addresses this problem by constraining the effective capacity of the model. In the present setting, regularization is not used merely to improve reconstruction quality for all samples. A useful reconstruction model should preserve contrast between dominant preceding-stage waveform patterns and newly emerged, less represented stage-specific deviations. A capacity-restricted autoencoder is therefore forced to allocate its limited representational resources to the most dominant and recurring patterns in the transition window. Signals that deviate from these dominant patterns receive less faithful reconstruction. In the context of adjacent-transition detection, when a subsequent damage stage introduces signal characteristics that differ from the preceding stage, the regularized model can produce a larger reconstruction gap between the two stages than an unrestricted model would.

The underlying principle is shared across regularization mechanisms, but each mechanism imposes the constraint differently. Dropout limits reliance on fragile unit-specific features through stochastic hidden-unit deactivation, weight decay encourages smoother and lower-complexity mappings by penalizing large weights, and sparsity restricts latent activation so that each signal is represented using a compact set of features. In all cases, the goal is to prevent the autoencoder from memorizing all stage-specific waveform variations equally well, thereby preserving the reconstruction contrast that makes adjacent transitions detectable [19]. The tradeoff is that excessive regularization can suppress useful structural information together with nuisance variation, collapsing reconstruction quality across all stages and eliminating detection contrast in the opposite direction. The regularization strength must therefore be set within an effective range.

In this work, we investigate three regularization methods, namely dropout, weight decay, and sparsity, and use the vanilla autoencoder (AE) without regularization as a benchmark. The vanilla AE minimizes only the reconstruction loss $\mathcal{L}_{\text{AE}} = \mathcal{L}_{\text{rec}}$, with no explicit capacity constraint. In the following subsections, we introduce each regularization method, describe how it constrains model capacity, and specify its key hyperparameter.

3.3.1 Dropout Regularization (AE + Dropout)

Dropout is an implicit regularization technique that restricts model capacity through stochastic training. During training, hidden units are randomly deactivated with probability p at each forward pass. In supervised learning, dropout is commonly interpreted as a method for improving generalization. In the present reconstruction-based transition-detection setting, its role is more specific. The objective is not to reconstruct every waveform equally well, but to preserve reconstruction contrast between dominant preceding-stage patterns and newly emerged stage-specific deviations. Dropout forces the autoencoder to rely on distributed, repeatable waveform structures rather than fragile sample-specific details. As a result, the learned reconstruction mapping becomes less capable of memorizing subtle subsequent-stage variations. At inference time, the preceding-stage patterns can still be reconstructed reliably, whereas newly introduced or less represented stage-specific components are reconstructed with lower fidelity, leading to an enlarged

PCC gap. The key parameter is the dropout rate p : higher values impose stronger capacity restriction but may degrade reconstruction quality if set too aggressively.

3.3.2 Weight Decay Regularization (AE + Weight Decay)

Weight decay adds an L_2 penalty on the trainable parameters to the reconstruction loss:

$$\mathcal{L}_{\text{wd}} = \mathcal{L}_{\text{rec}} + \lambda \|\Theta\|_2^2, \quad (7)$$

where Θ denotes the full set of trainable parameters and λ controls the penalty strength. This penalty discourages large parameter values, effectively shrinking the model toward smoother and lower-complexity reconstruction mappings. By limiting the magnitude of learned weights, weight decay reduces the model's ability to fit fine-grained sample-specific or stage-specific variations, including transient environmental disturbances. Its applicability is therefore strongest when the detection problem benefits from a globally smoother reconstruction function, although excessive shrinkage can reduce sensitivity to subtle but meaningful structural changes.

3.3.3 Sparsity Regularization (AE + Sparsity)

Sparsity regularization constrains the average activation of hidden units to remain close to a low target value ρ :

$$\mathcal{L}_{\text{sp}} = \mathcal{L}_{\text{rec}} + \beta \sum_j \text{KL}(\rho \parallel \hat{\rho}_j), \quad (8)$$

where $\hat{\rho}_j$ is the empirical average activation of hidden unit j across the training batch and $\text{KL}(\cdot \parallel \cdot)$ denotes the Kullback–Leibler divergence. The coefficient β controls the strength of the sparsity penalty. By encouraging most hidden units to remain inactive for any given input, sparsity regularization forces the autoencoder to represent each signal using only a compact subset of latent features. This selective activation restricts the model's representational flexibility and prevents it from encoding input-specific details that do not contribute to dominant reconstruction patterns. Sparsity is particularly suitable when only a limited subset of latent factors is expected to describe the recurring structural response, but an overly strong sparsity penalty may remove weak damage-sensitive components together with nuisance variation.

3.3.4 Hyperparameter Settings

The regularization hyperparameter values used in the main comparison are: dropout rate $p = 0.2$, weight decay coefficient $\lambda = 1 \times 10^{-5}$, and sparsity coefficient $\beta = 0.005$. For the sparsity model, the target activation ρ is fixed, and one-dimensional sensitivity analysis is conducted only over β . Sensitivity analysis for all three methods is reported in [Section 4](#).

3.4 Evaluation Protocol

To compare regularization strategies, we define pair coverage as the fraction of adjacent transitions that satisfy the detected-pair criterion:

$$\text{Coverage} = \frac{\#\{\text{detected adjacent pairs}\}}{12} \times 100\%. \quad (9)$$

The denominator of 12 corresponds to the total number of adjacent transitions in the 13-stage sequence ($D_1 \rightarrow D_2$ through $D_{12} \rightarrow D_{13}$). Coverage captures how broadly a given regularization method improves

detection across the full progression of structural stages, rather than focusing on isolated best-case performance. In addition, hyperparameter sensitivity is assessed through the mean AUROC computed over all 12 adjacent transitions as a function of the regularization strength parameter.

To support structured interpretation of the results, we define three representative case categories based on the detected-pair criterion and the regularization comparison:

1. **Detected:** the transition pair satisfies both the AUROC and PCC-gap thresholds.
2. **Difficult-I (weak contrast):** regularization may slightly improve separation, but the pair still does not meet the detected-pair criterion due to insufficient structural contrast between stages.
3. **Difficult-II (environment masking):** the pair does not support reliable structural interpretation because environmental fluctuations dominate the signal variation, obscuring any underlying structural change.

This categorization reflects deployment reality, where not all transitions are equally detectable and different failure modes require different mitigation strategies.

4 Results and Discussion

4.1 Pair Coverage Comparison

Table 3 summarizes the pair-level detection outcomes for all four model variants across the 12 adjacent transitions. The vanilla AE detects 2 of 12 pairs (16.7%), while dropout, weight decay, and sparsity detect 5 (41.7%), 6 (50.0%), and 7 (58.3%), respectively. All three regularized variants substantially outperform the unregularized baseline, confirming that regularization is a critical design variable for reconstruction-based detection under uncontrolled conditions [17]. As shown in Table 3, the additional pairs detected by the regularized models span a wide range of transitions, from early minor-damage stages such as $D_3 \rightarrow D_4$ to later severe-damage stages such as $D_{11} \rightarrow D_{12}$, indicating that regularization broadens detection coverage across the entire damage progression rather than benefiting only a narrow subset of transitions.

Table 3: Summary of pair-level detection outcomes and coverage comparison. Pair notation (i, j) denotes the transition $D_i \rightarrow D_j$.

Method	Regularization	Detected Pairs	Coverage (%)	Specific Pairs
AE	No	2/12	16.7	(4, 5), (7, 8)
AE + Dropout	Yes	5/12	41.7	(3, 4), (4, 5), (6, 7), (7, 8), (9, 10)
AE + Weight Decay	Yes	6/12	50.0	(3, 4), (4, 5), (6, 7), (7, 8), (9, 10), (11, 12)
AE + Sparsity	Yes	7/12	58.3	(3, 4), (4, 5), (6, 7), (7, 8), (8, 9), (9, 10), (11, 12)

To examine whether the detected-pair conclusion depends on a single operating point, Table 4 reports coverage under nearby AUROC and PCC-gap thresholds. An AUROC threshold of 0.80 is relatively permissive, whereas 0.90 is stricter for fine-grained adjacent-stage detection. Similarly, a PCC gap of 0.005 may admit weakly separated pairs, whereas 0.015 may exclude subtle but meaningful adjacent transitions. The selected criterion, $\text{AUROC} \geq 0.85$ and $\text{PCC gap} \geq 0.010$, provides a balanced operating point. Across nearby threshold choices, the regularized AE variants consistently outperform the unregularized AE, indicating that the main conclusion is not an artifact of a single threshold setting.

Table 4: Sensitivity of detected-pair coverage to AUROC and PCC-gap thresholds. Each entry reports the number of detected adjacent transitions out of 12. The selected operating point used in the main experiments is highlighted in bold.

AUROC Threshold	PCC-Gap Threshold	AE	AE+Dropout	AE+Weight Decay	AE+Sparsity
0.80	0.005	5/12	5/12	6/12	7/12
0.80	0.010	2/12	5/12	6/12	7/12
0.80	0.015	1/12	2/12	3/12	3/12
0.85	0.005	5/12	5/12	6/12	7/12
0.85	0.010	2/12	5/12	6/12	7/12
0.85	0.015	1/12	2/12	3/12	3/12
0.90	0.005	2/12	4/12	6/12	6/12
0.90	0.010	1/12	4/12	6/12	6/12
0.90	0.015	1/12	2/12	3/12	3/12

4.2 Analysis of Representative Detected Pairs

Figs. 2 and 3 present PCC trajectory plots for representative detected transitions, illustrating two distinct successful detection regimes.

Fig. 2 shows the transitions $D_4 \rightarrow D_5$ and $D_7 \rightarrow D_8$, which correspond to a moderate-contrast regime. In these cases, the vanilla AE already provides partial separation, but trajectory overlap and temporal instability limit detection confidence. The regularized variants increase separation clarity and temporal stability, making these transitions more reliably distinguishable. These observations indicate that regularization primarily sharpens an existing contrast.

Fig. 3 shows the transitions $D_6 \rightarrow D_7$ and $D_9 \rightarrow D_{10}$, which correspond to a baseline-ineffective regime. In these cases, the vanilla AE provides little useful discrimination between the two stages. The regularized variants, however, achieve clearly improved separability, demonstrating that regularization enables detection of transitions that the unregularized baseline cannot resolve.

4.3 Analysis of Representative Difficult Pairs

Fig. 4 presents PCC trajectory plots for two representative difficult transitions that remain undetected under all model variants. These cases illustrate two distinct failure regimes.

The left panel of Fig. 4 shows the transition $D_2 \rightarrow D_3$, which represents a weak-contrast regime. Regularization provides a slight improvement in separation, but the resulting PCC gap remains insufficient to satisfy the detected-pair criterion. This indicates that the structural contrast between the two stages is inherently small, limiting what reconstruction-based methods can resolve.

The right panel of Fig. 4 shows the transition $D_{10} \rightarrow D_{11}$, which represents an environment-dominated failure case. Unlike the weak-contrast case, the PCC trajectories in this transition are strongly disturbed during periods of low temperature and unusually high humidity. This observation is consistent with the original dataset validation, which reports that irregular environmental variations, including precipitation, snow, direct sunlight, and freezing conditions, can substantially distort guided waves and may produce changes comparable to or larger than damage-induced variations [5]. Therefore, the failure of this transition should not be interpreted solely as a limitation of the regularized AE. Rather, it reflects a practical deployment regime in which environmental masking reduces the reliability of reconstruction-based structural interpretation.

A full quantitative separation between environmental masking and structural progression is beyond the scope of this study and will require dedicated environmental compensation or weather-stratified modeling. Nevertheless, the representative case in Fig. 4 highlights why adjacent transition detection under uncontrolled long-term monitoring is more challenging than endpoint damage detection under laboratory conditions. Together, these two regimes clarify why certain transitions remain unresolved in Table 3: one is limited by weak structural contrast between stages, while the other is dominated by environmental interference that obscures the underlying structural change.

4.4 Hyperparameter Sensitivity

Fig. 5 presents the sensitivity analysis for each regularization scheme. The figure shows mean AUROC across all 12 adjacent transitions as a function of each regularization hyperparameter. The best mean AUROC values are 0.9071 for dropout, 0.9249 for weight decay, and 0.9288 for sparsity. This ranking is consistent with the pair coverage results in Table 3.




Figure 5: Hyperparameter sensitivity analysis using mean AUROC. Dashed lines indicate the selected regularization hyperparameter values: dropout rate = 0.2, weight decay coefficient = 1×10^{-5} , and sparsity coefficient = 0.005.

Three observations are relevant for practical tuning. First, all three regularization methods exhibit both effective and non-optimal intervals, indicating that hyperparameter tuning remains necessary even after the regularization method has been selected. Second, weight decay and sparsity demonstrate broader effective ranges than dropout under the current data setting, which is consistent with prior guided-wave regularization studies [18]. Third, the selected regularization hyperparameter values (dropout rate = 0.2, weight decay coefficient = 1×10^{-5} , sparsity coefficient = 0.005) provide stable defaults for this dataset and can serve as initial values for similar deployment scenarios.

4.5 Practical Implications

The experimental results support three practical implications. First, regularization consistently improves adjacent stage-transition detection across the full damage progression, confirming its importance as a design variable for unsupervised reconstruction-based monitoring under uncontrolled conditions. Second, under strong environmental interference, regularization provides less improvement than in moderate conditions, but still outperforms the unregularized baseline, suggesting that regularization alone cannot fully compensate for dominant environmental disturbances. Third, hyperparameter tuning remains necessary for all three regularization methods, as each exhibits both effective and non-optimal intervals in the sensitivity analysis.

The scope of the comparison should also be interpreted carefully. Because the proposed evaluation target is reconstruction-based adjacent-stage separability under a fixed local transition-window protocol, direct comparison with broader unsupervised SHM methods is not one-to-one. Methods based on explicit environmental compensation, stored baselines, reference-only training, latent likelihood, adversarial generation, or temporal state-space modeling rely on different assumptions and produce different detection scores. Applying them to the proposed 4.5-day + 0.5-day transition-window protocol would require nontrivial redesign of the input window, learning assumption, score function, and thresholding rule, introducing confounding factors beyond the regularization mechanism. Therefore, this study focuses on an intra-framework comparison under a fixed reconstruction pipeline that isolates the effect of regularization, while unified cross-method benchmarking is left for future work.

Detailed pair-level plots, full transition-wise AUROC/PCC-gap values, and path-wise coverage results are provided in the Supplementary Materials. Supplementary Table S1 lists all adjacent-transition metrics and shows that the bold detected entries reproduce the main coverage values in [Table 3](#). Supplementary Table S2 reports coverage across all eight guided-wave paths, showing that the benefit of regularization is not unique to path 6-1, although absolute coverage varies because of path-dependent sensor drift and environmental distortion.

5 Conclusions

This paper presents a controlled empirical study of how regularization improves unsupervised adjacent stage-transition detection for long-term guided-wave monitoring under uncontrolled conditions. Using one fixed autoencoder pipeline and a uniform detected-pair criterion ($\text{AUROC} \geq 0.85$ and average PCC gap ≥ 0.010), the vanilla AE detects 2 of 12 adjacent transitions, while dropout, weight decay, and sparsity detect 5, 6, and 7 transitions, respectively. The additional pairs detected by the regularized models span the full damage progression, from early minor-damage stages to later severe-damage stages, confirming that regularization broadens detection coverage rather than benefiting only a narrow subset of transitions.

The results also reveal the limitations of regularization. Under strong environmental interference, regularization provides less improvement than in moderate conditions, although it still outperforms the unregularized baseline. Hyperparameter sensitivity analysis shows that weight decay and sparsity exhibit broader effective ranges than dropout, but tuning remains necessary for all three methods. The selected hyperparameter values provide stable defaults for this dataset and can serve as starting points for similar deployment scenarios.

Because the main controlled analysis is based on a single-channel protocol (path 6-1), the primary conclusions should be interpreted as path-specific rather than as a claim of universal cross-system generalization. The supplementary path-wise analysis indicates that the regularization benefit is not unique to path 6-1, but cross-dataset validation, multi-channel fusion, and adjacent-transition uncertainty estimation remain important directions for future work.

Acknowledgement: Not applicable.

Funding Statement: The authors received no specific funding for this study.

Author Contributions: The authors confirm contribution to the paper as follows: study conception and design: Boyang Zhang and Yuang Geng; data collection: Boyang Zhang and Yuang Geng; analysis and interpretation of results: Boyang Zhang and Kang Gao; draft manuscript preparation: Boyang Zhang, Yuang Geng, Kang Gao, Jianan Gu and Li Ai. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The data that support the findings of this study are openly available in the figshare repository associated with the Scientific Data publication on long-term guided-wave monitoring (DOI: [10.6084/m9.figshare.28112504](https://doi.org/10.6084/m9.figshare.28112504)).

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

Supplementary Materials: The supplementary material is available online at <https://www.techscience.com/doi/10.32604/sdhm.2026.083281/sl>.

References

1. Figueiredo E, Park G, Farrar CR, Worden K, Figueiras J. Machine learning algorithms for damage detection under operational and environmental variability. *Struct Health Monit.* 2010;10(6):559–72. doi:10.1177/1475921710388971.
2. Moll J, Kathol J, Fritzen CP, Moix-Bonet M, Rennoch M, Koerdt M, et al. Open guided waves: online platform for ultrasonic guided wave measurements. *Struct Health Monit.* 2018;18(5–6):1903–14. doi:10.1177/1475921718817169.
3. Moll J, Kexel C, Kathol J, Fritzen CP, Moix-Bonet M, Willberg C, et al. Guided waves for damage detection in complex composite structures: the influence of omega stringer and different reference damage size. *Appl Sci.* 2020;10(9):3068. doi:10.3390/app10093068.
4. Kudela P, Radzienski M, Moix-Bonet M, Willberg C, Lugovtsova Y, Bulling J, et al. Dataset on full ultrasonic guided wavefield measurements of a CFRP plate with fully bonded and partially debonded omega stringer. *Data Brief.* 2022;42(5–6):108078. doi:10.1016/j.dib.2022.108078.
5. Kang Y, Yang Z, Yang H, Zhou J, Zhang ZR, Wang L, et al. Dataset on guided waves from long-term structural health monitoring under uncontrolled and dynamic conditions. *Sci Data.* 2025;12(1):1123. doi:10.1038/s41597-025-05300-5.
6. Croxford AJ, Wilcox PD, Drinkwater BW, Konstantinidis G. Strategies for guided-wave structural health monitoring. *Proc Royal Soc Math Phys Eng Sci.* 2007;463(2087):2961–81.
7. Su Z, Ye L. Identification of damage using Lamb waves: from fundamentals to applications. Cham, Switzerland: Springer Science & Business Media; 2009.
8. Croxford AJ, Moll J, Wilcox PD, Michaels JE. Efficient temperature compensation strategies for guided wave structural health monitoring. *Ultrasonics.* 2010;50(4–5):517–28. doi:10.1016/j.ultras.2009.11.002.
9. Harley JB, Moura JMF. Scale transform signal processing for optimal ultrasonic temperature compensation. *IEEE Trans Ultrason Ferroelect Freq Control.* 2012;59(10):22. doi:10.1109/tuffc.2012.2448.
10. Douglass ACS, Harley JB. Dynamic time warping temperature compensation for guided wave structural health monitoring. *IEEE Trans Ultrason Ferroelect Freq Control.* 2018;65(5):851–61. doi:10.1109/tuffc.2018.2813278.
11. Liu C, Harley JB, Berges M, Greve DW, Oppenheim IJ. Robust ultrasonic damage detection under complex environmental conditions using singular value decomposition. *Ultrasonics.* 2015;58(1):75–86. doi:10.1016/j.ultras.2014.12.005.
12. Wang Z, Cha YJ. Unsupervised deep learning approach using a deep auto-encoder with a one-class support vector machine to detect damage. *Struct Health Monit.* 2020;20(1):406–25. doi:10.1177/1475921720934051.
13. Nick W, Shelton J, Asamene K, Esterline AC. A study of supervised machine learning techniques for structural health monitoring. *Int J Mach Learn Comput.* 2015;1353:36. doi:10.7763/ijmlc.2015.v5.526.
14. Bull L, Worden K, Manson G, Dervilis N. Active learning for semi-supervised structural health monitoring. *J Sound Vib.* 2018;437(19):373–88. doi:10.1016/j.jsv.2018.08.040.
15. Sakurada M, Yairi T. Anomaly detection using autoencoders with nonlinear dimensionality reduction. In: *Proceedings of the MLSDA, 2014 2nd Workshop on Machine Learning for Sensory Data Analysis. MLSDA'14.* New York, NY, USA: ACM; 2014. p. 4–11. doi:10.1145/2689746.2689747.
16. Mao J, Wang H, Spencer BF. Toward data anomaly detection for automated structural health monitoring: exploiting generative adversarial nets and autoencoders. *Struct Health Monit.* 2020;20(4):1609–26. doi:10.1177/1475921720924601.

17. Gorgin R, Luo Y, Wu Z. Environmental and operational conditions effects on Lamb wave based structural health monitoring systems: a review. *Ultrasonics*. 2020;105(2087):106114. doi:10.1016/j.ultras.2020.106114.
18. Kang Y, Kim S, Harley JB. Unsupervised long-term damage detection in an uncontrolled environment through optimal autoencoder. *Mech Syst Signal Process*. 2023;199:110473. doi:10.1016/j.ymssp.2023.110473.
19. Kang Y, Liu T, Yang Z, Zhou Y, Tian Z, Kim NH, et al. Baseline optimized autoencoder-based unsupervised anomaly detection in uncontrolled dynamic structural health monitoring. *Struct Health Monit*. 2025;20(15):8204. doi:10.1177/14759217251324107.
20. Kang Y, Yang Z, Tian Z, Tetali H, Harley JB. Weight decay optimized unsupervised autoencoder based anomaly detection in uncontrolled dynamic structural health monitoring. In: *Dynamic data driven applications systems*. Cham, Switzerland: Springer Nature; 2026. p. 31–9. doi:10.1007/978-3-031-94895-4_3.
21. Fang C, Yang X, Gryllias K, Vandepitte D, Liu X, Zhang L, et al. Using removable sensors in structural health monitoring: a Bayesian methodology for attachment-to-attachment uncertainty quantification. *Mech Syst Signal Process*. 2025;224:111973. doi:10.1016/j.ymssp.2024.111973.
22. Ferreira PM, Lopez I, Sarfaraz R, Dervilis N, Worden K, Cross EJ, et al. Bayesian calibration for lamb wave propagation on a composite plate using a machine learning surrogate. *Mech Syst Signal Process*. 2024;208(7):111011. doi:10.1016/j.ymssp.2023.111011.
23. Fernandez-Navamuel A, Magalhaes F, Zamora-Sanchez D, Omella AJ, Garcia-Sanchez D, Pardo D. Deep learning enhanced principal component analysis for structural health monitoring. *Struct Health Monit*. 2022;21(4):1710–22. doi:10.1177/14759217211041684.
24. Li L, Yan J, Wen Q, Jin Y, Yang X. Learning robust deep state space for unsupervised anomaly detection in contaminated time-series. *IEEE Trans Knowl Data Eng*. 2023;35(6):6058–72. doi:10.1109/tkde.2022.3171562.
25. Eybpoosh M, Berges M, Noh HY. An energy-based sparse representation of ultrasonic guided-waves for online damage detection of pipelines under varying environmental and operational conditions. *Mech Syst Signal Process*. 2017;82(3–5):260–78. doi:10.1016/j.ymssp.2016.05.022.
26. Rautela M, Senthilnath J, Monaco E, Gopalakrishnan S. Delamination prediction in composite panels using unsupervised-feature learning methods with wavelet-enhanced guided wave representations. *Compos Struct*. 2022;291:115579. doi:10.1016/j.compstruct.2022.115579.
27. Lomazzi L, Junges R, Giglio M, Cadini F. Unsupervised data-driven method for damage localization using guided waves. *Mech Syst Signal Process*. 2024;208:111038. doi:10.1016/j.ymssp.2023.111038.
28. Silva MF, Santos A, Santos R, Figueiredo E, Costa JC. Damage-sensitive feature extraction with stacked autoencoders for unsupervised damage detection. *Struct Control Health Monit*. 2021;28(5):e2714. doi:10.1002/stc.2714.
29. Rastin Z, Ghodrati Amiri G, Darvishan E. Unsupervised structural damage detection technique based on a deep convolutional autoencoder. *Shock Vib*. 2021;2021(1):6658575. doi:10.1155/2021/6658575.
30. Ma X, Lin Y, Nie Z, Ma H. Structural damage identification based on unsupervised feature-extraction via variational auto-encoder. *Measurement*. 2020;160(1851):107811. doi:10.1016/j.measurement.2020.107811.