Vehicle Speed Invariant Drive-By Bridge Damage Detection

VARUN SAHAY, JATIN AGGARWAL, SEOYOUNG OH, JINGXIAO LIU and HAE YOUNG NOH

ABSTRACT

Given the central role of bridges in transportation networks, continuous monitoring of these structures is crucial to detect damage, ensure long-term serviceability, and prevent catastrophic failures. Traditional inspection methods, however, are costly, labor-intensive, and often subjective. Sensor-based approaches, such as strain gauges or accelerometers installed directly on bridges, require significant installation and maintenance efforts, limiting their scalability. An emerging alternative leverages vehicle–bridge interaction: by analyzing dynamic responses recorded by in-vehicle sensors, bridge conditions can be assessed without installing dedicated instrumentation on the structure. Vehicle accelerations reflect bridge-induced vibrations during crossings and indirectly encode dynamic properties of the bridge. However, variability in vehicle speed poses a significant challenge, as it affects the vibration signatures captured by accelerometers while remaining independent of bridge health. This confounding factor hinders the extraction of reliable, damage-sensitive features.

To overcome this challenge, we propose a vehicle speed invariant drive-by bridge damage detection model that employs adversarial learning to extract features from vehicle-induced vibrations that are sensitive to structural damage while invariant to vehicle speed. The model integrates long short-term memory (LSTM) layers with a Gradient Reversal Layer (GRL). The LSTM layers capture temporal patterns in the frequency components, enabling the extraction of features that reflect subtle temporal variations across the full vibration spectrum. The GRL imposes speed invariance by adversarially optimizing the feature representation to maximize accuracy in damage classification while minimizing performance in speed prediction. We evaluate our method through a lab-scale experimental vehicle–bridge system with the vehicle running at varying speeds. Our model performs $1.4\times$ better than baseline methods for bridge damage detection, achieving an accuracy of 91.01% and an F1-score of 93.29%.

Varun Sahay, Graduate Student, Email: varuns19@stanford.edu. Department of Civil and Environmental Engineering, Stanford University, Stanford, CA, USA

INTRODUCTION

Bridges are essential to both the economy and public safety, as they facilitate the efficient movement of goods and people, reducing travel times by connecting critical infrastructure. Yet, continuous deterioration over time makes them increasingly vulnerable to damage, posing significant structural and public safety risks. Early detection of such damage is essential to prevent catastrophic failures. Traditionally, bridge inspections are carried out by human inspectors, but these are typically limited to once every two years due to a shortage of skilled labor and the high cost of manual inspections. Bridge health monitoring through sensors directly installed on the bridge, such as accelerometers, strain gauges, geophones, and tiltometers, can provide detailed information on bridge health. However, installing and maintaining these sensors is costly and labor-intensive, making large-scale deployment impractical [1]. Further, each bridge of interest must be individually instrumented, which limits scalability. Recently, researchers have explored the use of vehicle-mounted accelerometers for bridge health monitoring by capturing vehicle vibrations when traversing the bridge. Vehicle accelerations capture bridge dynamic responses through vehicle–bridge interaction and thus encode bridge dynamics characteristics that are informative of the bridge's health condition [2–7]. This approach provides a scalable and cost-effective solution, as sensors mounted on vehicles can be deployed across multiple bridges. However, a key challenge remains: variations in vehicle speed can substantially affect the recorded acceleration responses [8], potentially leading to misclassification of bridge conditions.

To this end, we introduce a vehicle speed invariant bridge damage detection model that adversarially learns features from vehicle vibrations, which are sensitive to bridge structural damage yet invariant to vehicle speed. The method involves two components: a) capturing the temporal evolution of frequency-domain vibration amplitudes using long short-term memory (LSTM) [9] layers, and b) imposing the speed invariance in the extracted features through a Gradient Reversal Layer (GRL) [10, 11]. Firstly, vehicle-mounted accelerometers capture bridge-induced vibrations, whose frequency content varies over time and encodes information about bridge structural condition. Traditional learning approaches that rely on static or averaged frequency-domain features [12, 13], can miss critical low-amplitude, time-varying components that are highly sensitive to damage. To capture these dynamic signatures, we use LSTM layers to learn how vibration amplitudes at different frequencies evolve temporally, enabling the model to identify subtle patterns associated with damage. Secondly, vehicle speed also affects the recorded acceleration signals, making it difficult to disentangle speed-induced variability from actual structural changes. To address this, a GRL is employed before the speed prediction task. The GRL enables adversarial training, where the network is optimized to perform well on damage classification while intentionally performing poorly on speed prediction. This encourages the model to extract features that are sensitive to structural damage but invariant to vehicle speed. We evaluated the performance of our method on a lab-scale experimental bridge, with vehicle speeds ranging from 0.44 to 0.69 m/s. Our approach achieved an F1 score of 93.29% in damage detection. Furthermore, our method performs $1.4\times$ better than the baseline model (without the LSTM layers) in the bridge damage detection task.

VEHICLE SPEED INVARIANT BRIDGE DAMAGE DETECTION (VSI-BDD)

Our model (as shown in Figure 1) consists of two main components: (a) extracting temporal-frequency features from spectrogram inputs using LSTM layers, and (b) imposing speed invariant feature learning through a Gradient Reversal Layer (GRL).

First, to capture the temporal evolution of vehicle acceleration frequency amplitudes, we employ a recurrent neural network architecture with LSTM layers. Temporal variations in the frequency content of vehicle acceleration signals can reveal critical information about a bridge's structural condition, including early signs of damage. However, conventional models that rely on time-averaged representations often fail to capture these dynamic patterns—particularly within low-amplitude frequencies that may carry subtle yet informative damage signatures. To address this limitation, raw acceleration signals are segmented into 1-second windows to standardize input length and are then transformed into the time-frequency domain using the Short-Time Fourier Transform (STFT) spectrogram, which characterizes how frequency amplitudes evolve over time. These STFT spectrogram is then fed into a stack of LSTM layers. The output from the final time step of the last LSTM layer summarizes the sequence and is passed through a fully connected layer to extract features for both bridge damage classification and vehicle speed prediction.

Second, to ensure that the features extracted by the LSTM are invariant to vehicle speed, a GRL is introduced before the speed prediction task. The acceleration signals recorded from vehicles are inherently influenced by the vehicle's speed, introducing variability that can obscure the true structural response of the bridge. This speed-induced variability poses a significant challenge, as it can degrade the model's ability to accurately detect bridge damage. To overcome this, the GRL enables adversarial learning by acting as an identity function during the forward pass, while reversing the gradient (by multiplying it with $-\lambda$) during backpropagation. The negative sign causes the reversal, while λ controls how much the gradient from the speed prediction branch influences the feature extractor, pushing it to learn features that are less affected by speed. This process encourages the model to learn feature representations that minimize the damage classification loss while simultaneously maximizing the speed prediction loss, thereby making the extracted features invariant to speed and sensitive to damage.

EVALUATION

We evaluate the performance of our VSI-BDD model on the dataset obtained from experiments conducted on a lab-scale bridge-vehicle system. The bridge is 8 feet long and 2 feet wide, and the vehicle is equipped with two accelerometers—one at the front and one at the rear (Figure 2). Acceleration signals were recorded as the vehicle traversed the bridge at varying speeds, ranging from 0.44 to 0.69 m/s. The data were sampled at a frequency of 2000 Hz. Damage was introduced to the bridge by removing bolts at the midspan and quarter-span locations. For our detection model, both midspan (damaged type 1) and quarter-span (damaged type 2) damage states are grouped and labeled as *Damaged*.

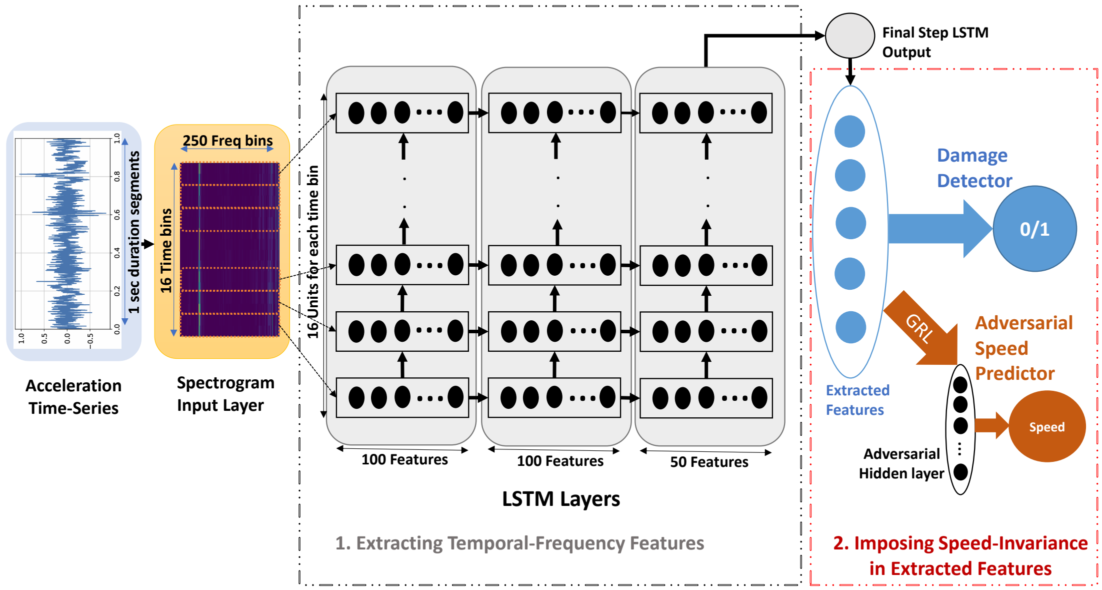

Figure 1. Architecture of the proposed vehicle speed invariant bridge damage detection (VSI-BDD) model. The network processes vehicle accelerometer signals via STFT, followed by a stack of LSTM layers to extract temporal features. These features are then passed to two branches: a damage classification branch and an adversarial speed prediction branch. A GRL is applied to the speed prediction branch to impose speed invariant feature learning.

Each acceleration signal had different lengths due to variability in vehicle speeds. To address these variations and ensure consistent model performance, we standardized the input data by segmenting it into uniform time windows of one second, corresponding to 2000 data points. Each segment was labeled as either *damaged* or *undamaged*, depending on the state of the bridge during data collection. The input to the LSTM model had the shape $(N, 16, 250)$, where N is the batch size and $(16, 250)$ is the STFT features extracted from 1-sec acceleration segments. A total of $19, 582$ samples were obtained, of which 90% were used for training and validation, while the remaining 10% were reserved for testing. The samples were shuffled to ensure a similar distribution across the training, validation, and test sets. For each sample, the corresponding output labels (0 for undamaged and 1 for damaged) and speed labels (continuous values computed by dividing the bridge length by the signal duration) were also obtained. Hyperparameters were tuned to optimize performance on the validation split (10% of the training set) using the Hyperband tuner [15], with dropout regularization applied to avoid overfitting. Specific details of the best-performing model are summarized in Table I.

Our model is evaluated on the lab-scale experiment with respect to a baseline model (Autoencoder+Classifier & Spectrogram+LSTM). The baseline Autoencoder+Classifier is a supervised learning model that performs dimensionality reduction of the segmented acceleration time-histories but lacks the ability to capture the temporal features. While

Aspect	Details
Optimizer	Adaptive Moment (Adam) [14]
Learning Rate	0.001
Loss Functions	- Damage Detection: Binary Crossentropy - Speed Prediction: Mean Squared Error (MSE)
Metrics	- Damage Detection: Accuracy, F1 Score - Speed Prediction: MSE
Activation Function	Rectified Linear Units (ReLU) except final outputs - Sigmoid for Damage Detection Output - Linear for the Velocity Predictor Output
Dropout Rate	0.1 after each LSTM layer
Feature Layer	Dense layer with 5 units
Adversarial weighing factor (λ)	0.01
Batch Size (N)	64

TABLE I. Details of the VSI-BDD model

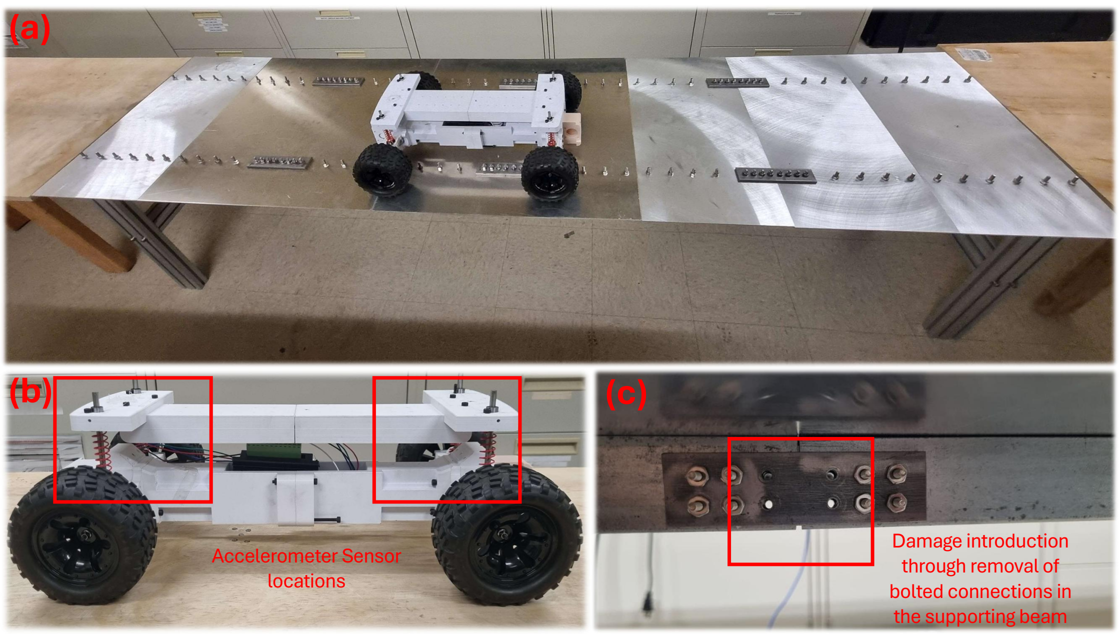

Figure 2. Lab-scale experimental setup (a) the steel girder with aluminum deck bridge (8ft x 2ft) with the vehicle model, (b) the vehicle model with the locations of the installed accelerometers, and (c) bridge damage simulated through removal of bolts from the gusset plate connection of longitudinal steel girder.

the Spectrogram + LSTM model enables us to capture the temporal features, but are not explicitly trained to learn vehicle speed invariant features. Our model (VSI-BDD) achieves an overall accuracy of 91.01% and an F1 score of 93.29% on the unseen test dataset. Table II shows the performance of our model compared to the baseline. Our model successfully outperforms the baseline and performs $1.4\times$ better than the Autoencoder+Classifier model.

Further, the invariance of vehicle speed is reinforced by the plot of each of the ex-

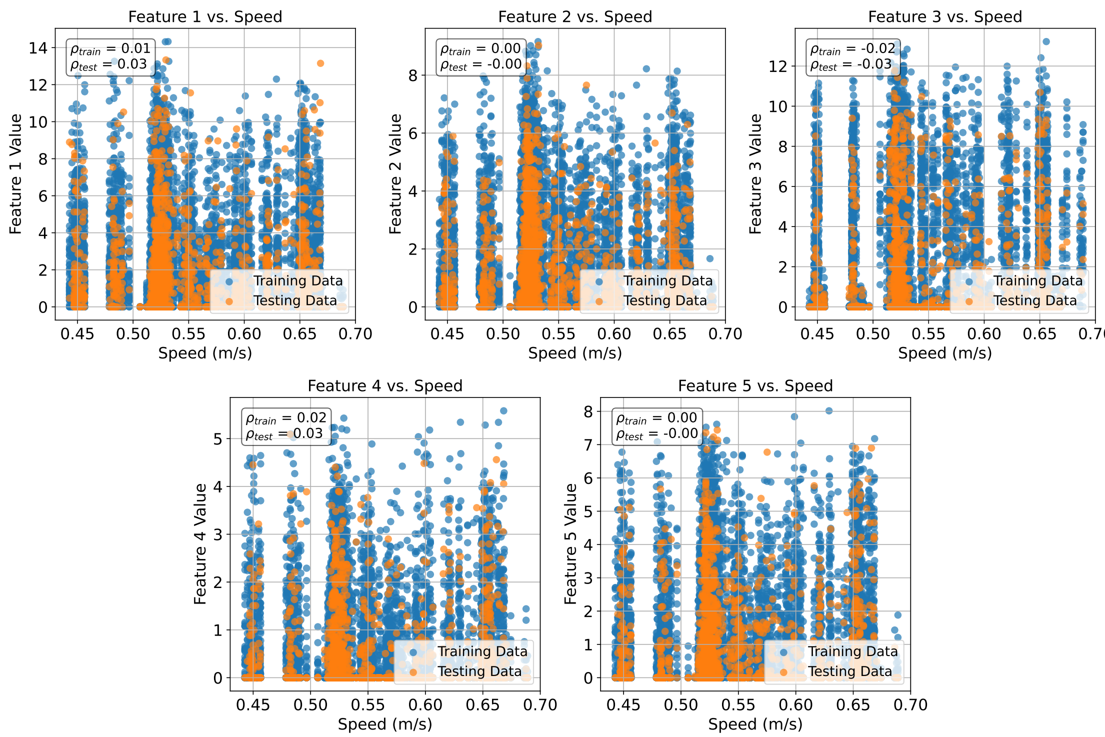

Figure 3. The scatter plots between the five extracted features and the vehicle speeds for both the training and testing datasets show that the Spearman's correlation coefficients are very close to zero for all five features. This implies that there is little to no monotonic relationship between vehicle speed and the values of the extracted features.

tracted features from the fully connected layer with vehicle speed (Figure 3). It is observed that the feature value does not show any correlation with vehicle speed for both the training and testing datasets. Figure 4 shows the t-SNE [16] plot of the extracted features for both damaged and undamaged cases, with the color of the scatter points indicate the damage state, showing a clear distinction. These signify the effectiveness of the model for the detection of bridge damage while being invariant to vehicle speed. Finally, Table III shows the results of our model across different speed ranges. It shows that the F1 score remains high across the entire range of speeds.

Model	**Accuracy**	**Recall**	**F1 Score**
VSI-BDD (Spectrogram + Adversarial LSTM)	91.01%	95.13%	93.29%
Spectrogram + LSTM	88.20%	88.55%	90.80%
Autoencoder + Classifier	65.73%	100%	78.98%

TABLE II. Performance of the our VSI-BDD model (Spectrogram + Adversarial LSTM) as compared to the baseline method on the damage detection while being speed invariant. Our model achieves overall F1-score of 93.29% in damage detection.

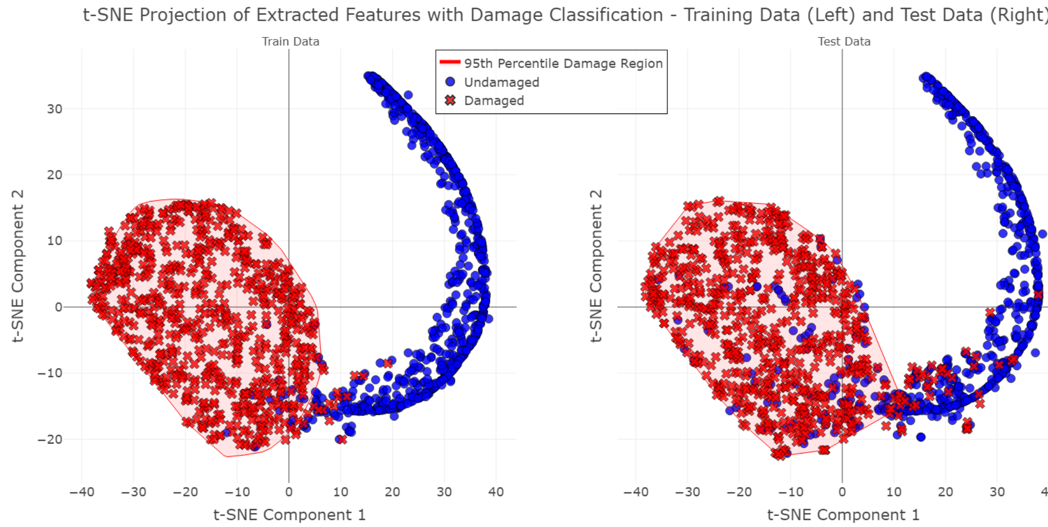

Figure 4. The t-SNE embedding plot of the extracted features for the damaged and undamaged bridge cases under varying vehicle speeds. The 95^{th} percentile damage region is shown by a red bounding box. A clear distinction of clusters between damaged and undamaged bridge cases is observed for both training data (left) and testing data (right).

No.	Speed Range (m/s)	Accuracy (%)	Recall (%)	F1 Score (%)
1	0.44-0.49	91.21	93.57	93.84
2	0.49-0.54	93.76	97.57	95.26
3	0.54-0.59	91.46	93.94	93.66
4	0.59-0.64	85.63	90.60	89.83
5	0.64-0.69	85.67	93.37	88.95

TABLE III. Performance metrics of our model across different speed ranges. The model performs good across all speed ranges showing the robustness of our method to different speeds.

CONCLUDING REMARKS

We present a vehicle speed invariant bridge damage detection method that leverages acceleration signals recorded during vehicle traversal to identify structural damage regardless of speed variations. By integrating a recurrent neural network with a Gradient Reversal Layer, the proposed model extracts features that are both sensitive to bridge condition and invariant to changes in vehicle speed. The approach achieves strong performance, with an F1-score of 93.29%, outperforming the baseline by a factor of $1.4\times$. These results demonstrate the potential of vehicle-mounted sensors for accurate and scalable bridge health monitoring under normal traffic conditions. Future work may extend this framework to classify damage types or localize defects, while preserving speed invariance.

REFERENCES

1. Aggarwal, J., J. Liu, and H. Noh. 2023. “Vehicle-invariant Drive-by Monitoring Across Multiple Bridges through Bootstrapping enhanced Unsupervised Domain Adaptation,” in *Structural Health Monitoring*.
2. Lederman, G., Z. Wang, J. Bielak, H. Noh, J. Garrett, S. Chen, J. Kovačević, F. Cerda, and P. Rizzo. 2014. “Damage quantification and localization algorithms for indirect SHM of bridges,” in *Bridge Maintenance, Safety, Management and Life Extension - Proceedings of the 7th International Conference of Bridge Maintenance, Safety and Management, IABMAS 2014*, Taylor and Francis - Balkema, Bridge Maintenance, Safety, Management and Life Extension - Proceedings of the 7th International Conference of Bridge Maintenance, Safety and Management, IABMAS 2014, ISBN 9781138001039, pp. 640–647, doi:10.1201/b17063-93, 7th International Conference on Bridge Maintenance, Safety and Management, IABMAS 2014 ; Conference date: 07-07-2014 Through 11-07-2014.
3. Lederman, G., S. Chen, J. Garrett, J. Kovačević, H. Y. Noh, and J. Bielak. 2017. “Track-monitoring from the dynamic response of an operational train,” *Mechanical Systems and Signal Processing*, 87:1–16, ISSN 0888-3270, doi: https://doi.org/10.1016/j.ymssp.2016.06.041.
4. Lederman, G., S. Chen, J. H. Garrett, J. Kovačević, H. Y. Noh, and J. Bielak. 2017. “Track monitoring from the dynamic response of a passing train: A sparse approach,” *Mechanical Systems and Signal Processing*, 90:141–153, ISSN 0888-3270, doi: https://doi.org/10.1016/j.ymssp.2016.12.009.
5. Lederman, G., S. Chen, J. H. Garrett, J. Kovačević, H. Y. Noh, and J. Bielak. 2017. “A data fusion approach for track monitoring from multiple in-service trains,” *Mechanical Systems and Signal Processing*, 95:363–379, ISSN 0888-3270, doi: https://doi.org/10.1016/j.ymssp.2017.03.023.
6. Liu, J. et al. 2020. “Diagnosis algorithms for indirect structural health monitoring of a bridge model via dimensionality reduction,” *Mechanical Systems and Signal Processing*, 136:106454, doi:10.1016/j.ymssp.2020.106454.
7. Liu, J. et al. 2023. “HierMUD: Hierarchical multi-task unsupervised domain adaptation between bridges for drive-by damage diagnosis,” *Structural Health Monitoring*, 22(3):1941–1968, doi:10.1177/14759217221081159.
8. Yang, Y. B. and C. W. Lin. 2005. “Vehicle–bridge interaction dynamics and potential applications,” *Journal of Sound and Vibration*, 284(1-2):205–226, doi:10.1016/j.jsv.2004.06.032.
9. Hochreiter, S. and J. Schmidhuber. 1997. “Long Short-Term Memory,” *Neural Computation*, 9(8):1735–1780.
10. Ganin, Y., E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky. 2016. “Domain-Adversarial Training of Neural Networks,” *Journal of Machine Learning Research*, 17(59):1–35.
11. Ganin, Y. and V. Lempitsky. 2015. “Unsupervised Domain Adaptation by Backpropagation,” in *International Conference on Machine Learning (ICML)*, pp. 1180–1189.
12. Kustiana, W. A. A., B. R. Trilaksono, M. Riyansyah, S. A. Putra, W. Caesarendra, G. Królczyk, and M. Sulowicz. 2024. “Bridge Damage Detection with Support Vector Machine in Accelerometer-Based Wireless Sensor Network,” *Journal of Vibration Engineering & Technologies*, 12(1):21–40, ISSN 2523-3939, doi:10.1007/s42417-024-01400-5.
13. Bao, Y., C. Song, W. Wang, T. Ye, L. Wang, and L. Yu. 2013. “Damage Detection of Bridge Structure Based on SVM,” *Mathematical Problems in Engineering*, 2013(1):490372, doi: https://doi.org/10.1155/2013/490372.
14. Kingma, D. P. and J. Ba. 2014. “Adam: A Method for Stochastic Optimization,” *arXiv*

preprint arXiv:1412.6980.

15. Li, L., K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar. 2018. “Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization,” *Journal of Machine Learning Research*, 18(185):1–52.
16. van der Maaten, L. and G. Hinton. 2008. “Visualizing Data using t-SNE,” *Journal of Machine Learning Research*, 9(Nov):2579–2605.