Reconstruction of 2D Guided Wave Propagation Data Using Generative Adversarial Networks

SZYMON DZIEWIC, MICHAL DZIENDZIKOWSKI
and ZIEMOWIT DWORAKOWSKI

ABSTRACT

Guided wave propagation analysis is one of the critical areas in Structural Health Monitoring (SHM) that enables the detection and assessment of damage in critical infrastructure. Visualization of wave interactions with structural anomalies, such as holes or cracks, gives important insights into the integrity and behavior of materials. In particular, automatic decision systems taking guided wave data as input can be used to monitor vital infrastructure, provided that they are trained on large amounts of diverse data.

Unfortunately, due to the cost, complexity, and time-consuming nature of experimental measurements, obtaining high-quality labeled data for guided wave propagation analysis remains a significant challenge. These restrictions on generating diverse datasets often make it more difficult to train reliable machine learning models for SHM, especially when handling a variety of complex damage scenarios.

Generative Adversarial Network (GAN) can offer a solution to those issues by producing realistic synthetic datasets that replicates key features of experimentally obtain data, allowing for cost-effective augmentation of SHM datasets. Those generative models are well suited for this task, based on their ability to simulate accurate representations of guided wave propagation through damaged structures, capturing subtle interactions between waves and structural anomalies. This allows for augmentation of SHM datasets at a reasonable cost, which eventually improves machine learning models ability to generalize damage scenarios that haven't been seen and reducing dependency on expensive experimental setups.

In this paper we demonstrate how to recreate two-dimensional (2D) guided wave propagation data through the aluminum plate with damages modeled as flat-bottom holes. Our method produces high-fidelity images with very close resemblance to training data obtained using laser vibrometer. This approach highlights the potential of generative models

Szymon Dziewic, AGH University of Krakow, Poland
Michał Dziendzikowski, Air Force Institute of Technology, Poland
Warsaw University of Technology, Poland Ziemowit Dworakowski, AGH University of Krakow, Poland

to improve data diversity and availability for SHM applications, in particular in simulation of previously unseen damage cases which enable directed augmentation of existing SHM databases. This strategy can reduce the gap between limited experimental datasets and machine learning-driven SHM systems.

INTRODUCTION

Guided-wave SHM can detect damage through elastic waves but often it relies on extensive labeled data, which are costly to collect [1]. GAN [2] have emerged in computer vision and, more recently, in SHM, to synthetically expand training sets. However, unconditional GANs lack fine control over temporal offsets and sensor configurations. For example, in [3], a conventional GAN was used to synthesize ultrasonic B-scans for flat-bottom and side-drilled holes, achieving high normalized cross-correlation and low Euclidean distance against real scans . In [4] StyleGAN2-inspired GAN was used on full-field guided-wave data from the OpenGuidedWaves database, demonstrating high-fidelity reconstruction of unseen waveforms and enabling damage classifiers trained solely on synthetic data . In [5] a ACGANs was used to produce realistic vibration signals for fault diagnosis on rotating machinery improving classifier performance under data scarcity, while in [6] 1-D WDCGAN-GP was used to generate labeled acceleration data that effectively balance imbalanced datasets .

In this paper, we propose a conditional Wasserstein GAN with gradient penalty (CWGAN-GP) tailored for 2D guided-wave future-frame prediction. Given a current wavefield frame and a normalized time offset, our U-Net generator [7] predicts the future frame, while a CNN discriminator assesses pair authenticity. We reinforce physical accuracy via a composite loss: Wasserstein adversarial [8, 9], masked L1, SSIM [10], CNN perceptual [11], and frequency-domain terms. We validate two training variants—First Run (L1+SSIM) and Second Run (feature loss, data augmentation, reflect-padding)—on laser-vibrometer data from a 12-sensor aluminum plate. Trainig for the Second Run started with best weight from the First Run. Results on validation sequences shows structural fidelity and generalization to unseen sensor configurations, offering a way to augmenting SHM datasets with realistic guided-wave sequences.

DATA AND EXPERIMENTAL SETUP

To train our conditional WGAN on realistic guided-wave propagation, we acquired full-field measurements on a 10 mm aluminum plate instrumented with 12 piezoelectric transducers in a 3×4 grid. Each transducer acted in turn as an actuator while a scanning laser vibrometer captured the out-of-plane response, yielding twelve sequences from

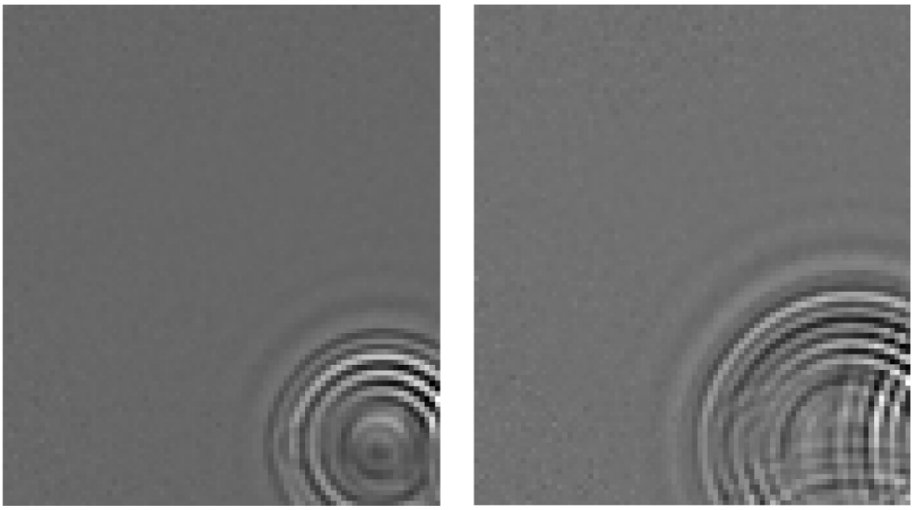

Figure 1. (left) wave at actuation; (right) subsequent propagation and reflections.

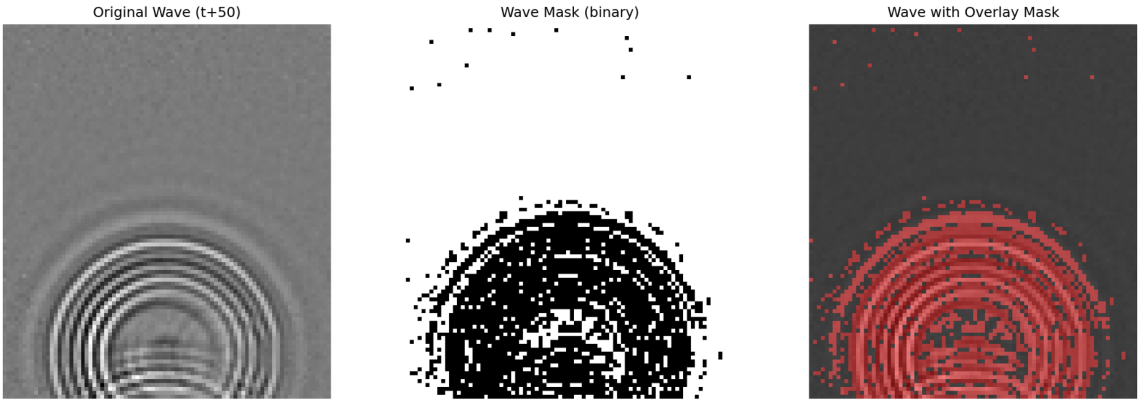

Figure 2. Binary masks highlighting active wavefront regions for loss weighting.

which we cut relevant 200 time-step wavefield frames that include Lamb-wave propagation and reflections (Fig. 1). Nine sequences were used for training and three for validation. For supervised, time-conditioned learning, each "current" frame was paired with multiple "future" frames at varying time offsets, normalized to [0,1], producing on the order of 190000 training examples. To emphasize dynamic wave regions, we computed binary masks of the active wavefront by thresholding amplitude and applied these masks in our loss weighting (Fig. 2). This approach ensures the model focuses on the key wave interactions needed for accurate future-frame prediction.

CWGAN-GP MODEL AND TRAINING

We implement a conditional WGAN-GP with a U-Net generator G and a CNN discriminator D. G receives the current wavefield frame plus a normalized time-offset channel and predicts the future frame, while D scores real versus generated pairs under the Wasserstein GAN with gradient penalty framework [8, 9, 12].

To ensure both stability and physical accuracy, we combine the adversarial loss with four complementary terms:

- **Masked L1 Loss:** weights errors on active wavefronts via binary masks.
- **SSIM Loss:** promotes structural similarity in the masked regions [10].
- **Feature Loss:** (Second Run only) captures higher-level wave patterns via a CNN perceptual metric [11].
- **Frequency Loss:** aligns the power spectrum of real and generated frames.

The generator minimizes the weighted sum:

$$L_G = L_{\text{adv}} + 40\, L_1 + 90\, L_{\text{SSIM}} + \gamma\, L_{\text{feat}} + 40\, L_{\text{freq}},$$

with $\gamma = 0$ for First Run and $\gamma = 60$ for Second Run. The discriminator loss is

$$L^{(D)} = \underbrace{\mathbb{E}_{\text{fake}}[D(x)] - \mathbb{E}_{\text{real}}[D(x)]}_{\text{WGAN}} + \lambda\, \mathbb{E}_{\hat{x}}\big(\|\nabla_{\hat{x}} D(\hat{x})\|_2 - 1\big)^2,$$

using $\lambda = 10$ (Run 1) or 7 (Run 2).

Training is carried out in TensorFlow/Keras on an NVIDIA RTX4080 with mixed precision for up to 20 epochs, using Adam Optimizer. We apply early stopping based on masked L1 + SSIM validation loss. Second Run adds reflect padding and simple augmentation (horizontal flip + Gaussian noise) to improve generalization. Table I summarizes the two configurations.

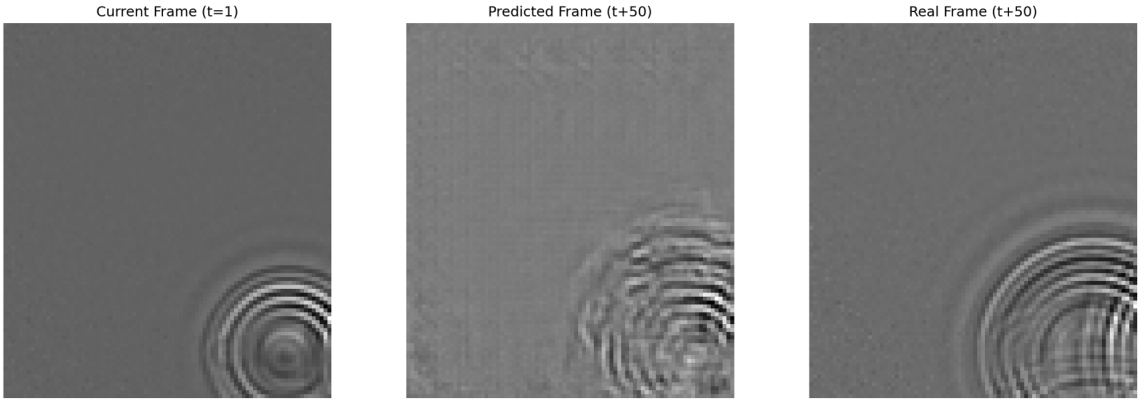

Figure 3. First Run Epoch 20 comparison: (left) current frame as and input for generator, (center) predicted image after 50 frames, (right) real image after 50 frames

TABLE I. Training Run Configurations (First vs Second Run)

Aspect	**First Run**	**Second Run**
Learning Rate	3×10^{-5}	1×10^{-5} (for stability)
D:G update ratio	7:1	4:1 (fewer D steps needed)
Gradient Penalty λ	10.0	7.0 (slightly lower)
Generator padding	Zero padding	Reflect padding
Feature Loss γ	0 (disabled)	60 (enabled)
Data Augmentation	None	Horizontal flip + noise
Wave Mask Threshold	0.09 (9% of max)	0.08 (8% of max)

RESULTS

We evaluated performance on three held-out sensor sequences, focusing on both visual fidelity and quantitative metrics. The First Run (Figure 3) initially showed blurry predictions that sharpen by Epoch 20 showing that generator identified wave patterns. The Second Run (Figure 4) due to starting from the best weights from First Run begins with more coherent patterns and preserves wave intersections with small blurring by Epoch 20. Some background speckle remains.

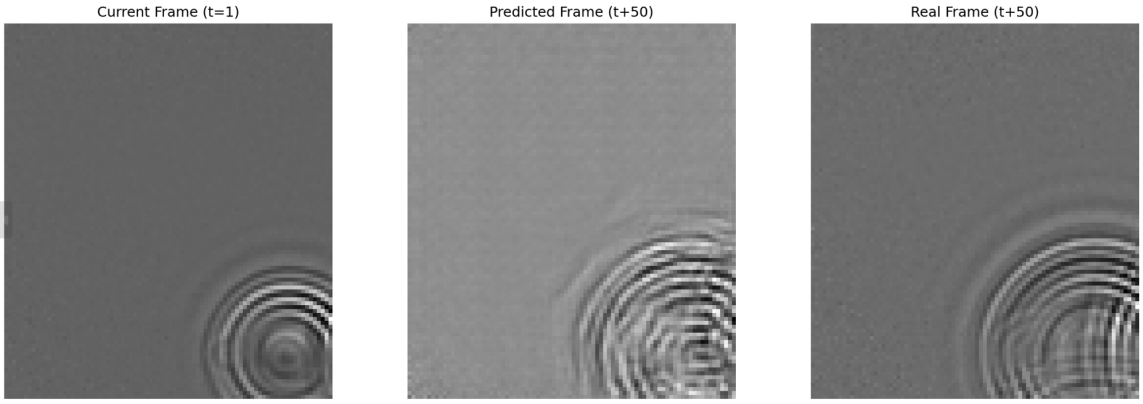

Figure 4. Second Run Epoch 20 comparison: (left) current frame as and input for generator, (center) predicted image after 50 frames, (right) real image after 50 frames

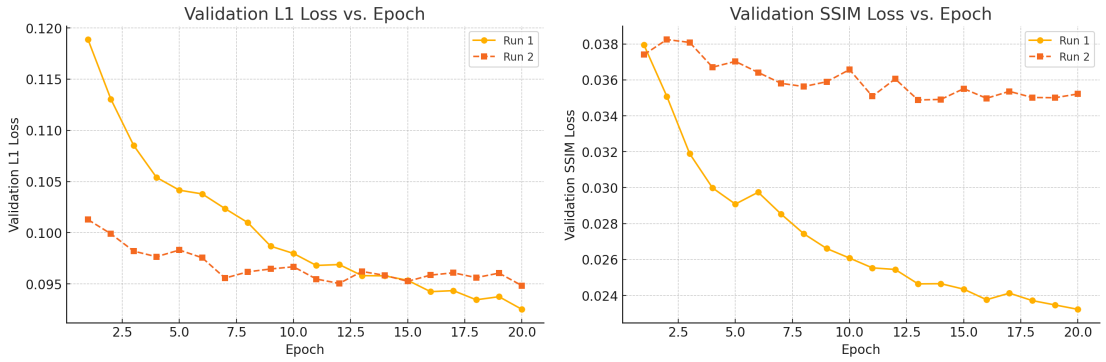

Figure 5. Validation metrics: (left) L1 loss and (right) SSIM loss across epochs. Second Run (orange) converges faster in L1 but plateaus in SSIM, while First Run (yellow) improves steadily.

Quantitatively, Figure 5 shows the validation metrics across epochs. In First Run (yellow), the average L1 loss drops steadily to around 0.093, while SSIM loss decreases to around 0.023. Second Run (orange) that started from loaded best weights from the First Run begins with a lower L1 and thanks to feature/frequency guidance converges quickly to 0.095. It's SSIM loss fluctuates around 0.036, slightly higher than First Run final 0.023, reflecting compromises when balancing multiple loss terms.

When tracking the combined losses (Figure 6), First Run loss falls from 9.41 to 6.94, while the second run starts at 7.42 and descends to 6.96. Since different validations were

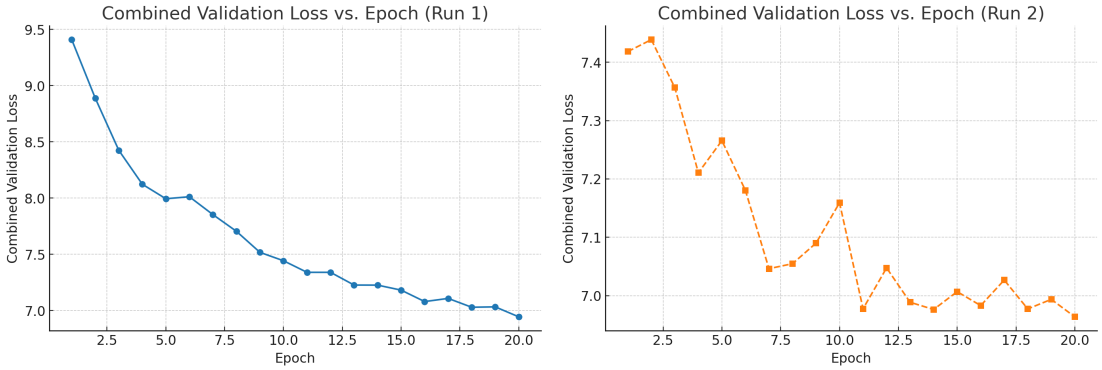

Figure 6. Combined validation loss vs. epoch: (left) First Run uses $40L_1 + 90(1 - \text{SSIM})$; (right) Second Run includes $60L_{\text{feature}}$.

used in each run, we cannot compare the values by itself. In the Second Run, we see the value fluctuating around 7, suggesting a performance ceiling under the current architecture and loss weights. Overall, our CWGAN-GP showed satisfactory structural integrity and simulation of future wave frames. Future improvements are indicated by residual background artifacts, under-prediction of edge-reflected amplitudes and amplitude decay in long-offset predictions. These include expanded edge-case augmentation, multi-step or recurrent conditioning for long-horizon forecasts, and more aggressive masking.

ACKNOWLEDGMENT

The work presented in this paper was supported by the National Centre for Research and Development in Poland, under project number LIDER13/0132/2022.

REFERENCES

[1] Farrar, C. R. and K. Worden. 2007. "An introduction to Structural Health Monitoring," *The Royal Society – Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 365:303–315.

[2] Goodfellow, I., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. 2014. "Generative Adversarial Networks," *Advances in Neural Information Processing Systems*, 27.

[3] Virupakshappa, K. and E. Oruklu. 2020. "Using Generative Adversarial Networks to Generate Ultrasonic Signals," in *Proceedings of the 2020 IEEE International Ultrasonics Symposium (IUS)*, IEEE, pp. 1–4.

[4] Heesch, M., M. Dziendzikowski, K. Mendrok, and Z. Dworakowski. 2022. "Diagnostic-Quality Guided Wave Signals Synthesized Using Generative Adversarial Neural Networks," *Sensors*, 22(3848), doi:10.3390/s22103848.

[5] Shao, S., P. Wang, and R. Yan. 2019. "Generative Adversarial Networks for Data Augmentation in Machine Fault Diagnosis," *Computers in Industry*, 106:85–93.

[6] Luleci, F., F. N. Catbas, and O. Avci. 2023. "Generative Adversarial Networks for Labeled Acceleration Data Augmentation for Structural Damage Detection," *Journal of Civil Structural Health Monitoring*, 13:181–198.

[7] Ronneberger, O., P. Fischer, and T. Brox. 2015. "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, Springer International Publishing, pp. 234–241.

[8] Arjovsky, M., S. Chintala, and L. Bottou. 2017. "Wasserstein GAN," *arXiv preprint arXiv:1701.07875*.

[9] Gulrajani, I., F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville. 2017. "Improved Training of Wasserstein GANs," in *Advances in Neural Information Processing Systems*, Curran Associates, Inc., vol. 30.

[10] Wang, Z., A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. 2004. "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, 13(4):600–612.

[11] Johnson, J., A. Alahi, and L. Fei-Fei. 2016. "Perceptual Losses for Real-Time Style Transfer and Super-Resolution," in *European Conference on Computer Vision*, Springer International Publishing, pp. 694–711.

[12] David, F. 2023. "Generative Deep Learning: Teaching Machines to Paint, Write, Compose, and Play," *O'Reilly*.