Available online at www.eccomasproceedia.org
Eccomas Proceedia SMART (2025) 828-839

ECCOMAS
Proceedia

XI ECCOMAS Thematic Conference on Smart Structures and Materials - SMART 2025
M. Krommer, M. Nader and M. Schagerl (Eds)

Received: 30 April 2025 / Revised: 4 August 2025 / Accepted: 25 October 2025
© The Author(s) 2025

Localization of Impacts on Thin-Walled Structures by Recurrent Neural Networks: End-to-End Learning from Real-World Data

Alexander Humer, Lukas Grasböck, Ayech Benjeddou

Abstract. Today, machine learning is ubiquitous, and structural health monitoring (SHM) is no exception. Specifically, we address the problem of impact localization on shell-like structures, where knowledge of impact locations aids in assessing structural integrity. Impacts on thin-walled structures excite Lamb waves, which can be measured with piezoelectric sensors. Their dispersive characteristics make it difficult to detect and localize impacts by conventional methods. In the present contribution, we explore the localization of impacts using neural networks. In particular, we propose to use recurrent neural networks (RNNs) to estimate impact positions end-to-end, i.e., directly from sequential sensor data. We deal with comparatively long sequences of thousands of samples, since high sampling rate are needed to accurately capture elastic waves. For this reason, the proposed approach builds upon Gated Recurrent Units (GRUs), which are less prone to vanishing gradients as compared to conventional RNNs. Quality and quantity of data are crucial when training neural networks. Often, synthetic data is used, which inevitably introduces a reality gap. Here, by contrast, we train our networks using physical data from experiments, which requires automation to handle the large number of experiments needed. For this purpose, a robot is used to drop steel balls onto an aluminum plate equipped with piezoceramic sensors. Our results show remarkable accuracy in estimating impact positions, even with a comparatively small dataset.

Key words: Structural Health Monitoring, Impact Localization, Guided Wave Propagation, Recurrent Neural Networks, Robot-Supported Experiments

Alexander Humer
Institute of Technical Mechanics, Johannes Kepler University Linz, Austria, e-mail: alexander.humer@jku.at

Lukas Grasböck
Linz Center of Mechatronics GmbH, Linz, Austria, e-mail: lukas.grasboeck@lcm.at

Ayech Benjeddou
Institut Supérieur de Mécanique de Paris, Saint Ouen, France & Roberval Laboratory, Université de Technologie de Compiègne (UTC), France, e-mail: ayech.benjeddou@isae-supmeca.fr

ISSN:3078-543X. Published by Eccomas Proceedia.
Peer-review under responsibility of the organizing committee of SMART 2025.
doi: 10.7712/150125.13682.208007

1 INTRODUCTION

Impact localization on thin-walled structures plays a crucial role in structural health monitoring (SHM) and damage assessment, particularly in aerospace and automotive applications where lightweight and safety-critical components are common. Knowledge of impact locations is essential for assessing structural integrity and detecting potential damage in structures. While this problem has traditionally been approached using blends of physics and data-based methods, recent advances in machine learning have opened new possibilities for more effective solutions, see, e.g., [1].

Mechanical impacts on thin-walled structures generate guided elastic waves—commonly referred to as *Lamb waves* [2], which can be captured using a network of piezoelectric sensors. Accurately determining the location of an impact in such structures, however, remains a challenging task due to the complex propagation characteristics of these waves. Lamb waves are highly dispersive and their behavior is further complicated in non-isotropic materials such as fiber-reinforced polymers, where anisotropy and inhomogeneity introduce additional complexity in wave propagation, see, e.g., [3, 4, 5]. Conventional approaches to impact localization predominantly rely on estimating the time-of-arrival (ToA) of the impact-induced wavefront at several distributed sensors. Classical methods from the field of signal processing include, among several others, continuous wavelet transform, modified energy-ratio and the Akaike information criterion, see, e.g., [6, 7] for an overview. These estimates are then used in triangulation-type algorithms to infer the impact location, see, e.g., [8]. While this methodology has been widely adopted, it typically utilizes only a small portion of the available signal information—namely, the wavefront arrival. The rich signal content that follows, including reflections at structural boundaries and mode conversions, is usually disregarded, despite its potential to enhance localization performance.

Approaches based on different kinds of neural networks can leverage a larger part of the information contained in sensor signals to accurately estimate impact positions. Recent contributions combine fully-connected networks for localization and recurrent structures to reconstruct force histories [9]. Spatio-temporal graph convolutional networks for the localization of impacts on composite plates have been proposed in [10]. The framework of Gaussian processes is used in [11] to account for the inherently probabilistic nature of impact-induced waves in thin-walled structures.

In this work, we propose an alternative approach that aims to exploit the full informational content of the sensor signals by learning to estimate the impact position directly from raw time-series data. In particular, we propose to use *recurrent neural networks* (RNNs) to estimate impact positions *end-to-end*, i.e., directly from *sequential sensor data*. Our central hypothesis is that an end-to-end learning framework can yield more accurate and robust localization results, potentially with fewer sensors than those required in conventional methods. Due to the high sampling rate needed to accurately capture Lamb waves, we deal with comparatively long sequences, which pose a significant challenge when training RNNs. Therefore, we employ more complex network architectures such as Gated Recurrent Units (GRUs) [12], which incorporate

gating mechanisms to manage long-term dependencies more efficiently than standard RNNs.

Quality and quantity of data are crucial when training neural networks. Often, synthetic data is used, which inevitably introduces a reality gap. Here, by contrast, we train our networks using physical data from experiments, which requires automation to handle the large number of experiments needed. For this purpose, a collaborative robot is used to drop steel balls onto an aluminum plate equipped with piezoceramic sensors.

Our results show remarkable accuracy in estimating impact positions, even with a comparatively small dataset. This suggests that the proposed approach effectively leverages the rich information content in Lamb wave signals, potentially offering a more efficient and accurate alternative to conventional impact localization methods.

The present paper is organized as follows: Following this introduction, we briefly outline the network architecture of the proposed estimator. Section 3 introduces the experimental setup for the fully automated generation of training data. The (minimal) data preperation and the training setup are described in Section 4. Results are presented and discussed in Section 5, highlighting the model's performance across different sensor configurations and signal durations. Finally, Section 6 summarizes our findings and outlines directions for future work.

2 METHODOLOGY

2.1 Impact location and sensor signals

The impacts on the plate occur in a target region $T \subset \mathbb{R}^2$. The vector of true positions with respect to some fixed Cartesian frame is denoted by $\boldsymbol{y} = (p_x, p_y)^T \in T$. Let S denote the number of sensors distributed across the structure to be monitored. We assume a fixed sampling rate, i.e., sensors take measurements at equidistant time points t_i, where $i = 1, \ldots, L$. At each time t_i, sensors record signals $x_i^{(j)} \in D$, for $j = 1, \ldots, S$, with D being the range of sensor signals, which are then combined into a vector x_i:

$$\boldsymbol{x}_i = \boldsymbol{x}(t_i) = \left(x^{(1)}(t_i), \ldots, x^{(S)}(t_i)\right)^T \in D \times \ldots \times D \subset \mathbb{R}^S. \tag{1}$$

For notational convenience, sensor signals $\boldsymbol{x}_i$ form the rows of a matrix $\mathbf{X}$,

$$\mathbf{X} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_L)^T \in \mathbb{R}^{L \times S}, \tag{2}$$

which represents the complete sensor data over time.

2.2 Model architecture

The proposed location estimator LE is supposed to predict the impact location $\boldsymbol{y}$ from the (temporal) sequence of sensor signals $\mathbf{X}$. In other words, LE maps the (spatio-)temporal patterns in the sensor data to a position in the target region T, which, from a mathematical point of view, translates into

$$LE : \mathbb{R}^{L \times S} \to \mathbb{R}^2 : \quad \mathbf{X} \mapsto \hat{\boldsymbol{y}} = LE(\mathbf{X}), \tag{3}$$

where a hat is introduced for predicted positions. Given the sequential nature of the sensor data $\mathbf{X}$, it is crucial to capture temporal dependencies for accurate predictions. RNNs are a natural choice to represent such an estimator, since, owing to their sequential nature, they are capable of processing variable-length input sequences while maintaining a memory of prior time steps. The memory of an RNN is represented by the hidden state $\boldsymbol{h}_i$, which is computed from the previous hidden state $\boldsymbol{h}_{i-1}$ and the current input $\boldsymbol{x}_i$,

$$\boldsymbol{h}_i = r(\boldsymbol{h}_{i-1}, \boldsymbol{z}_i) \in \mathbb{R}^H, \qquad \boldsymbol{z}_i = \mathbf{W}^{(\mathrm{in})}\boldsymbol{x}_i + \boldsymbol{b}^{(\mathrm{in})}, \qquad \mathbf{W}^{(\mathrm{in})} \in \mathbb{R}^{H\times S},\ \boldsymbol{z}_i, \boldsymbol{b}^{(\mathrm{in})} \in \mathbb{R}^H, \quad (4)$$

where H denotes the dimensionality of the hidden state. Note that we have introduced a linear layer to project the input vector $\boldsymbol{x}_i$ onto a vector $\boldsymbol{z}_i$, which has the same dimension as the hidden state. Following machine learning conventions, $\mathbf{W}^{(\mathrm{in})}$ is referred to as weight-matrix and $\boldsymbol{b}^{(\mathrm{in})}$ denotes the bias vector.

Intuitively, the prediction for the impact location is made once the whole sequence of sensor data has been processed, in our case, by a linear map applied to the final hidden state $\boldsymbol{h}_L$:

$$\hat{\boldsymbol{y}} = \mathbf{W}^{(\mathrm{out})}\boldsymbol{h}_L + \boldsymbol{b}^{(\mathrm{out})}, \qquad \mathbf{W}^{(\mathrm{out})} \in \mathbb{R}^{2\times H}, \boldsymbol{b}^{(\mathrm{out})} \in \mathbb{R}^2. \quad (5)$$

The specific RNN architecture is subject to practical constraints that arise from the particular SHM application of impact monitoring. Lamb waves propagate at comparatively high velocities. For the typical frequencies excited by impacts, the phase velocities of the fundamental anti-symmetric (A_0) mode are 1000 m s^{-1} or higher for the aluminum panel used in our experimental study; see, e.g., [7]. For this reason, high sampling rates are required to capture the propagation of Lamb waves, which implies long sequences of sensor data.

The long input sequences resulting from high sampling rates challenge standard RNNs, which are known to suffer from vanishing gradients when modeling long-term dependencies. GRUs, in contrast, introduce gating mechanisms that control the flow of information to mitigate the problem of vanishing gradients, while maintaining a lower complexity than Long Short-Term Memory (LSTM) [13].

2.3 Loss function and training

The estimator is trained in a supervised manner by minimizing the mean squared error (MSE) between predicted positions $\hat{\boldsymbol{y}}$ and true positions $\boldsymbol{y}$, which serves as natural choice for the loss function. Our initial experiments, however, have shown that the estimator is strongly conditioned to the length of the input sequence if the prediction is made only after the entire sequence has been processed. When training is performed with a fixed sequence length L, the model learns to make accurate predictions only after exactly L time steps, but neither for longer nor for shorter sequences. To address this issue, we modify the training procedure such that the estimator makes intermediate predictions at each time step t_i. By minimizing the loss not only at the end of the sequence, but for every time step, i.e.,

$$\mathrm{MSE} = \frac{1}{L}\sum_{i=1}^{L} \left\|\hat{\boldsymbol{y}}_i - \boldsymbol{y}\right\|^2, \quad (6)$$

the model learns to generalize across variable-length inputs. Using the above loss, we force the estimator to make predictions for every sensor data $\boldsymbol{x}_i$ being processed, so that the length of the training sequence is no longer "learned". This strategy is particularly natural in online monitoring systems, which are meant to provide location estimates continuously, i.e., sample by sample, as new data arrives. That raises the question, of course, what the model predicts prior to the arrival of impact-induced waves at the sensors, i.e., before information on the location of their source becomes is contained in the sensor signals. Initially, the prediction defaults to a mean estimation which reflects the distribution of impact locations in the training data, e.g., the center of the plate if impacts are uniformly distributed.

3 EXPERIMENTAL SETUP

As with any data-driven approach, both the quality and quantity of data are crucial for creating accurate and reliable estimators. Due to the limited availability of real-world data, it has become common practice—particularly in the AI domain—to use synthetic data, especially when training neural networks, which are notorious for their high requirements in terms of data. In physics-related problems, numerical simulation provides effective means to generate synthetic training data, with real-world experiments typically requiring significantly greater expenses. In the present application, however, the situation is different. Owing to the fine temporal and spatial resolution required, simulations of transient Lamb waves that propagate in thin-walled structures are prohibitively demanding in terms of computational efforts to create a reasonable amount of training data. For this reason, physical experiments are key to the realization of the proposed impact localization estimator. As a side effect, we avoid issues related to what is referred to as *reality gap*, i.e., an inevitable modeling error in simulation models.

In the present work, a nearly square aluminum plate measuring $x \times y \times z = 904\,\mathrm{mm} \times 902\,\mathrm{mm} \times 2\,\mathrm{mm}$ serves as an example of a thin-walled structure to be impacted. To capture the propagation of impact-induced waves, eight piezoelectric patch transducers (PI Ceramic P-876.SP1) denoted by S_i, with $i = 1, \ldots, 8$, are bonded to the plate's top surface in a distance of 100 mm from the edges, see Fig. 1 (a). Sensors S_1 – S_4 are placed at the corners of the plate; S_5 – S_8 are located at the centers of the plate's edges. The center positions of the sensors with respect to the global Cartesian frame placed in the upper left corner are listed in Tab. 1. The active area of the patches is 10 mm × 10 mm in size. We want to avoid impacts on or directly

Table 1: Positions of the sensors on the aluminum plate. The coordinates are referring to the centers of the sensor patches.

Coordinates	S_1	S_2	S_3	S_4	S_5	S_6	S_7	S_8
x / mm	100	803	100	803	452	100	803	452
y / mm	100	100	802	802	100	451	451	802

next to the sensors, which is why we leave a margin of 5 mm. The crosshatched region marks the area of potential impacts on the plate. The sensor signals are sampled at 1 MHz using a

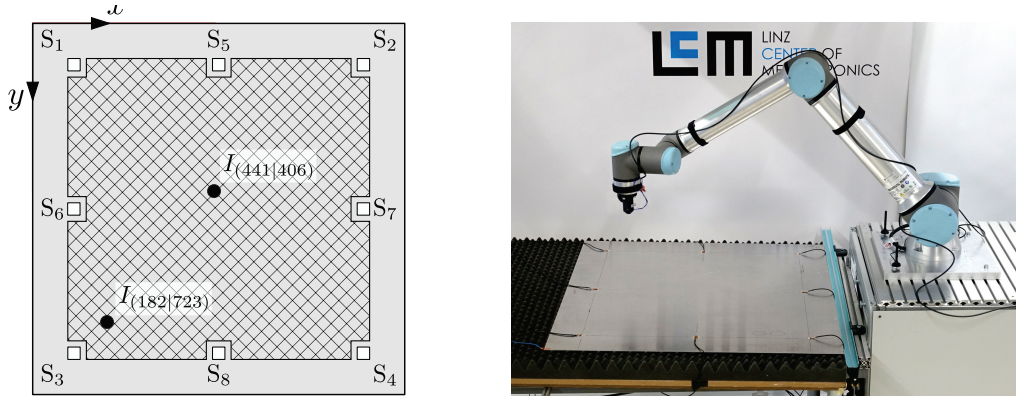


Figure 1: Experimental setup for impact localization on an aluminum plate: Schematic representation of the plate with the position of the sensors S_j, $j = 1, \ldots, 8$ (a). The crosshatched area indicates the region where impacts are applied. Two example impact positions from the dataset of 5000 impacts are marked as black dots, with their (x, y)-coordinates given in millimeters in parentheses. Picture of the robotic impact test setup used to generate the dataset for training neural networks (b).

National Instruments NI USB-6366 measurement system, which digitizes voltages in the range of ±10 V through a 16-bit analog-to-digital converter, corresponding to an amplitude resolution of 0.31 mV. To protect the data acquisition system from high voltages, a 1:10 probe is used.

To generate a sufficiently large and diverse dataset for training the estimator, a large number of impact events must be recorded. Manual generation of impacts is impractical and prone to inaccuracies in positioning. Therefore, a robotic system is employed to automate the experiments, ensuring precise and repeatable impact locations and enabling efficient collection of high-quality data as needed for training the estimator. For this purpose, a collaborative robot (Universal Robots UR10) is employed, see Fig. 1 (b). An end effector equipped with an electromagnet picks up a steel ball from a designated position and releases it at the precise impact location on the plate. The plate is supported from the bottom by a slightly inclined foam mat (less than 1°), ensuring that, after impact, the ball rolls into a guide rail that returns it to the pickup position, where a magnetic proximity sensor verifies the ball's return before initiating the next drop. The robot operates through a real-time data exchange (RTDE) interface [14].

4 DATA PREPARATION AND TRAINING

Within the present work, balls are dropped from a fixed height of 200 mm onto 5000 random positions sampled from a uniform distribution. Figure 2 shows two representative impacts and the corresponding signals of the corner sensors $\mathrm{S}_1 - \mathrm{S}_4$. The locations of the impacts are shown in Fig. 1. For a rather centered impact (441 mm, 406 mm), sensor signals are relatively similar from a qualitative point of view. The waves excited by the impact arrive at almost the same time

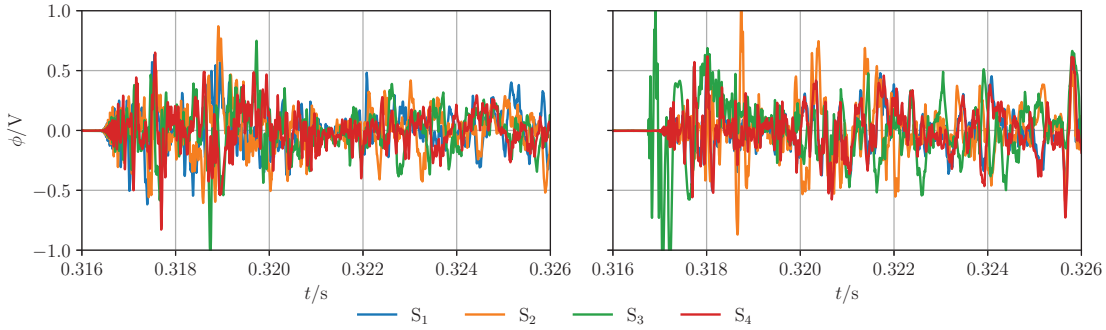


Figure 2: Representative sensor signals: Impacts at a centered position (441 mm, 406 mm) (left) and close to a corner (182 mm, 723 mm) (right).

at the four sensors. For an impact close to a corner (182 mm, 723 mm), the individual sensor signals differ significantly from one another. Unsurprisingly, the sensor closest to the impact (S_3) records waves earlier than the other three sensors. We also note a higher amplitude, since, due to the shorter distance, the waves attenuate less. The peak voltages ϕ are mostly in the range of ±1 V; the maximum voltage occurring in the dataset is 2.21 V. The noise level is three orders of magnitude smaller; we observe constant offsets in the sensor signals of up to 3 mV. Except for a downsampling to 250 kHz, no further processing is applied to the sensor signals.

Measurements are triggered by the robot, which turns off the magnet to drop the steel ball as soon as it reaches the desired position. Between the trigger and the impact on the plate, we see a consistent delay of approximately 316 ms. For efficient training, we do not want to include the entire non-informative pre-impact data. To avoid conditioning the network to a fixed arrival time, we apply a random cropping strategy around the detected signal onset. For each impact, we determine the time t_{th} at which the signal level rises beyond a threshold of 50 mV. We then crop a window of length T from the sensor signals, i.e., $t \in [t_{\mathrm{th}} - t_{\mathrm{shift}}, t_{\mathrm{th}} - t_{\mathrm{shift}} + T]$, where t_{shift} is randomly chosen from the interval $t_{\mathrm{shift}} \in [0.5\,\mathrm{ms}, 2\,\mathrm{ms}]$.

The target positions to be predicted by the estimator are provided in m, which implies that numerical values range from 0.1 to 0.8. No further normalization is applied in what follows. The dataset comprises 4000 impacts for training, 500 for validation (hyperparameter tuning, early stopping) and 500 for testing purposes. The proposed model is implemented in PyTorch [15] and trained on an NVIDIA A30 GPU.

5 RESULTS

The following results aim to demonstrate the feasibility of the proposed end-to-end learning approach for impact localization, rather than to achieve optimal performance through extensive hyperparameter tuning. To this end, we use a baseline model configuration based on empirical estimates that have proven effective in our training experiments. Specifically, the estimator consists of a single GRU layer with an input and hidden size of $H = 32$ ($H = 64$ when using eight sensors). Training is conducted over 5000 episodes using the AdamW optimizer with an

initial learning rate of 1×10^{-2} and a step-wise learning rate decay by a factor of 0.9 every 100 episodes. A fixed batch size of 500 is (mostly) used throughout training; batches are sampled randomly from the training set. Depending on the length T of the crop window, which ranges from 5 ms to 20 ms, the length of the input sequence varies from 1250 to 5000 samples, which is a considerable length in the context of RNNs.

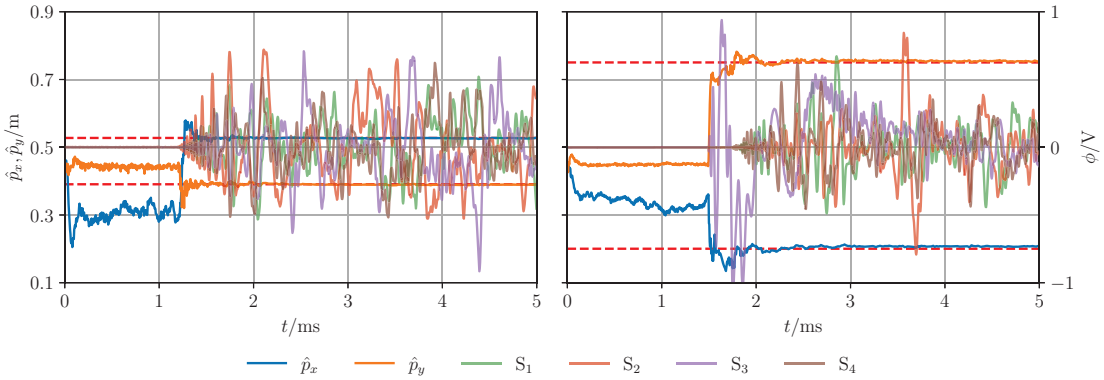


Figure 3: Sequential predictions of impact locations ($S = 4$, $T = 20\,\mathrm{ms}$) for impacts at (527 mm, 391 mm) (left) and (200 mm, 750 mm) (right): Predicted positions $\hat{p}_x$, $\hat{p}_y$ are illustrated along with corresponding input sensor signals $\mathrm{S}_1, \ldots, \mathrm{S}_4$. As long as sensor signals do not contain information on impacts, no meaningful estimates can be made. As impact-induced waves arrive at the sensors, estimated impact locations converge to (and remain at) true positions (red dashed lines).

Figure 3 illustrates the evolution of predicted impact locations $\hat{\boldsymbol{y}}$ over 5 ms as the sensor signals are sequentially processed for two positions of the test set, i.e., (527 mm, 391 mm) (left) and (200 mm, 750 mm) (right), using an estimator trained on $T = 20\,\mathrm{ms}$ sequences recorded by four sensors $\mathrm{S}_1, \ldots, \mathrm{S}_4$. Initially, the sensor signals do not contain any information related to an impact, which is why the predictions fluctuate around seemingly random values. As soon as impact-induced waves arrive at the sensor, the predictions quickly converge to the true positions, which are indicated by red dashed lines, and remain almost constant subsequently. For the second, off-center impact (right), we can observe that the estimator learns to predict impact locations using a single sensor only. Although a sensor signal is only present in S_3 at first, the estimated impact position is already close to the true position before the remaining sensors begin to measure.

Being computed over the entire input sequence, the loss is not an appropriate measure of the estimator's prediction quality. As an intuitive performance metric, we therefore compute the Euclidean distance between the true impact location $\boldsymbol{y}$ and the predicted location $\hat{\boldsymbol{y}}$, which is

averaged over the last $N = 10$ time steps:

$$e = \|\langle \hat{\boldsymbol{y}} \rangle - \boldsymbol{y}\|, \qquad \langle \hat{\boldsymbol{y}} \rangle = \frac{1}{N} \sum_{i=L-N}^{L} \hat{\boldsymbol{y}}_i. \tag{7}$$

Figure 4 illustrates the accuracy of our proposed estimator across all impacts in the test set. The true positions of impacts are represented by semi-transparent round markers, while the estimated positions are denoted by crosses. The color coding indicates the Euclidean distance between true and estimated positions, cf. Eq. (7), providing a visual representation of the distribution of the estimation error over the plate structure. For the sake of clarity, the upper limit of the colorbar is set to a value of 25 mm rather than the maximum error in order to emphasize nuances in the variation of the prediction quality. The results are organized in a table where rows correspond to different numbers and configurations of sensors used for prediction, while columns represent different time window lengths, i.e., $T = 5, 10, 20\,\mathrm{ms}$. Specifically, rows in Fig. 4 correspond to: two adjacent sensors $\mathrm{S}_1, \mathrm{S}_2$ in the top row (a)–(c), the four corner sensors S_1 – S_4 in the middle row (d)–(f), and all eight sensors S_1 – S_8 in the bottom row (g)–(i).

When using all eight sensors, the dimensionality of the input per time step increases accordingly. Therefore, the input and hidden size of the GRU cell are both increased to $H = 64$ to account for the larger amount of data that is being processed. For the longest time window ($T = 20\,\mathrm{ms}$), the batch size then has to be reduced for GPU memory constraints; for short windows ($T = 5\,\mathrm{ms}$), we increase the batch size to improve convergence and mitigate overfitting.

Table 2: Impact localization accuracy: Mean error and standard deviation of the position error e for different lengths of the input signals and different numbers of sensors (Note: **bold** values indicate variations from the base model architecture and training setup).

	Sensors	T/ms	L	H	B	mean(e)/mm	std(e)/mm
(a)	1 – 2	5	1250	32	500	9.146	5.845
(b)	1 – 2	10	2500	32	500	7.808	5.445
(c)	1 – 2	20	5000	32	500	6.417	4.945
(d)	1 – 4	5	1250	32	500	8.029	5.263
(e)	1 – 4	10	2500	32	500	6.110	4.941
(f)	1 – 4	20	5000	32	500	5.842	3.941
(g)	1 – 8	5	1250	**64**	**1000**	8.110	5.385
(h)	1 – 8	10	2500	**64**	500	5.902	3.633
(i)	1 – 8	20	5000	**64**	**250**	5.808	3.497

The mean prediction error and its standard deviation over the test set are listed in Tab. 2. Our results reveal a clear and consistent trend: longer input windows generally lead to better localization accuracy. This suggests that the model benefits from later parts of the sensor signals,

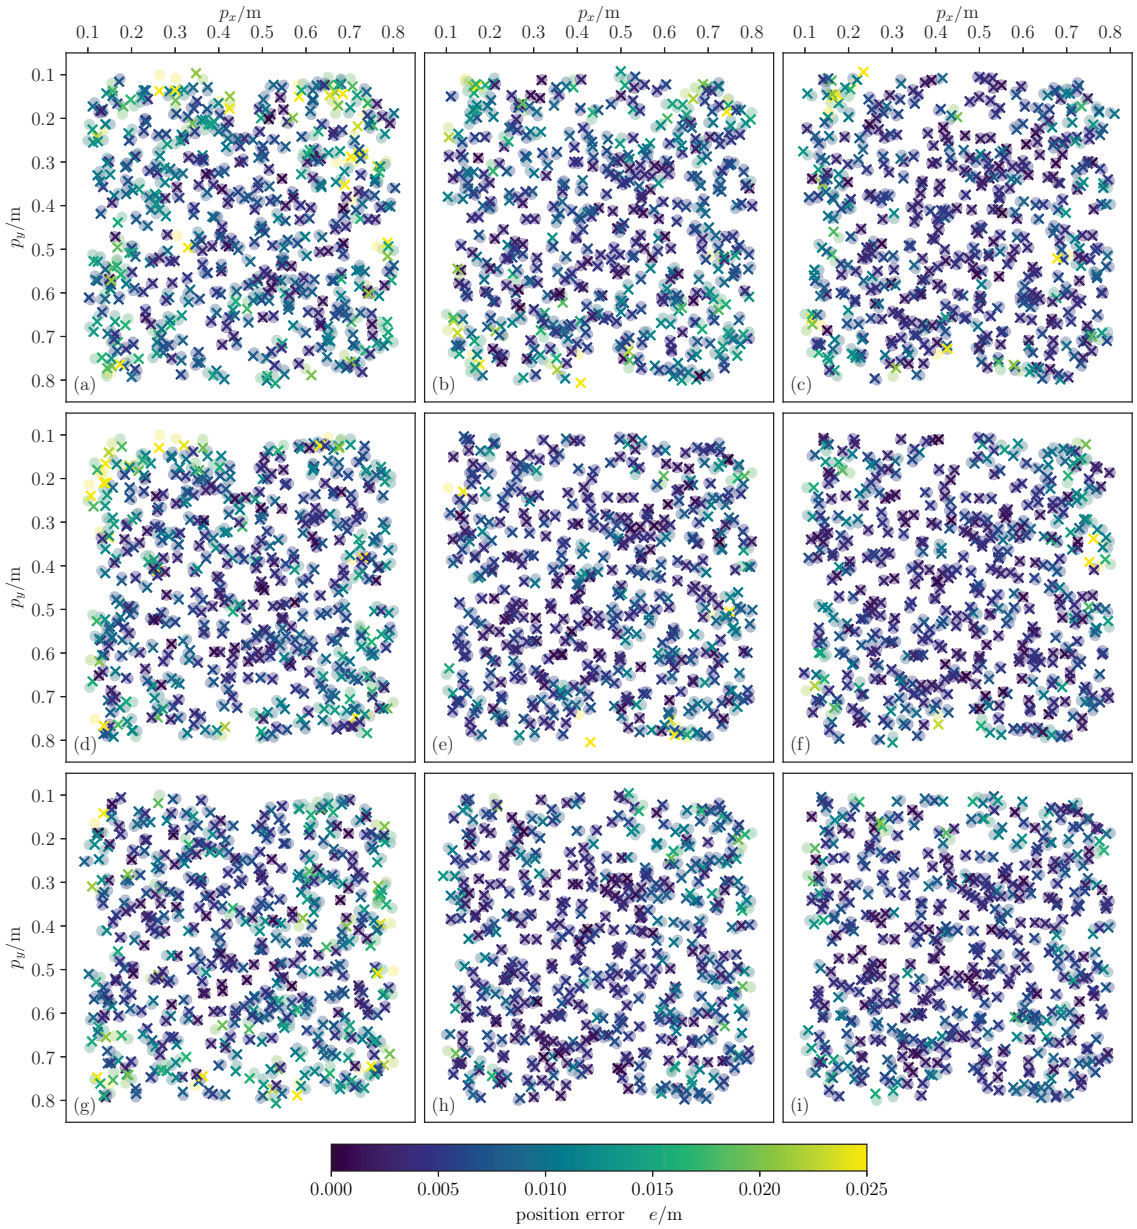


Figure 4: Impact localization accuracy (test set): True positions are indicated by semi-transparent round markers, estimated positions by crosses. Colors respresent the Euclidean distance between true and estimated positions. Rows correspond to the number of sensors used for the prediction ($S = 2, 4, 8$); columns represent different lengths of time windows ($T = 5, 10, 20\,\mathrm{ms}$).

such as reflections and features related to the dispersive nature of Lamb waves, rather than relying solely on the initial wave front. Additionally, using more sensors improves performance, although the gains from increasing the number of sensors from four to eight are relatively modest. Interestingly, even with just two sensors, the estimator is capable of delivering accurate predictions, which is something traditional triangulation approaches typically cannot achieve, since they usually require three sensors at least. We emphasize that the mean position errors are smaller than the dimensions of the piezoelectric patches.

6 CONCLUSION AND OUTLOOK

We presented an end-to-end approach for localizing impacts on thin-walled structures using GRU-based recurrent neural networks. Unlike conventional methods that rely on time-of-arrival estimation, our model learns directly from raw sensor signals and captures both early and late-arriving wave components. Trained entirely on experimental data collected through a robot-assisted setup, the method achieves high localization accuracy—even with a limited dataset and as few as two sensors.

Future work will focus on extending the method to more complex structural scenarios, including composite shells and reinforced structures. We also plan to investigate optimal sensor placement strategies and sensor shapes. Further, we want to explore more advanced network architectures, such as transformers, to further enhance performance of end-to-end impact localization.

REFERENCES

[1] L. Vu-Quoc and A. Humer, "Deep Learning Applied to Computational Mechanics: A Comprehensive Review, State of the Art, and the Classics," *CMES-Computer Modeling in Engineering and Sciences*, vol. 137, no. 2, pp. 1069–1343, 2023.

[2] H. Lamb, "On waves in an elastic plate," *Proceedings of the Royal Society of London. Series A, Containing papers of a mathematical and physical character*, vol. 93, no. 648, pp. 114–128, 1917.

[3] F. Ciampa, M. Meo, and E. Barbieri, "Impact localization in composite structures of arbitrary cross section," *Structural Health Monitoring*, vol. 11, no. 6, pp. 643–655, 2012.

[4] B.-W. Jang, Y.-G. Lee, C.-G. Kim, and C.-Y. Park, "Impact source localization for composite structures under external dynamic loading condition," *Advanced Composite Materials*, vol. 24, no. 4, pp. 359–374, 2015.

[5] D. Xiao, Z. Sharif-Khodaei, and M. Aliabadi, "Hybrid physics-based and data-driven impact localisation for composite laminates," *International Journal of Mechanical Sciences*, vol. 274, p. 109222, 2024.

[6] L. Grasboeck, A. Humer, and A. Benjeddou, "Revisiting multidisciplinary time-of-arrival estimation methods for impact detection on plates," in *Proceedings of the 10th ECCO-*

MAS thematic conference on smart structures and materials, D. Saravanos, A. Benjeddou, N. Chrysochoidis, and T. Theodosiou, Eds., 2023, pp. 107–118.

[7] L. Grasboeck, A. Humer, and A. Benjeddou, "Detection and time-of-arrival estimation of impact-induced waves in composite laminates," *Structural Health Monitoring*, vol. 24, no. 4, pp. 2428–2446, 2025. [Online]. Available: https://doi.org/10.1177/14759217241260848

[8] T. Kundu, "Acoustic source localization," *Ultrasonics*, vol. 54, no. 1, pp. 25–38, 2014.

[9] C. Huang, W. Liao, H. Sun, Y. Wang, and X. Qing, "A hybrid FCN-BiGRU with transfer learning for low-velocity impact identification on aircraft structure," *Smart Materials and Structures*, vol. 32, no. 5, p. 055012, apr 2023.

[10] Z. Zhao and N.-Z. Chen, "Spatial-temporal graph convolutional networks (stgcn) based method for localizing acoustic emission sources in composite panels," *Composite Structures*, vol. 323, p. 117496, 2023.

[11] D. Xiao, Z. Sharif-Khodaei, and M. H. Aliabadi, "Robust impact localisation on composite aerostructures using kernel design and bayesian fusion under environmental and operational uncertainties," 2025. [Online]. Available: https://arxiv.org/abs/2501.18393

[12] K. Cho, B. van Merrienboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1724–1734.

[13] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[14] Universal Robots, "Real-Time Data Exchange (RTDE) Guide," 2024.

[15] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "PyTorch: an imperative style, high-performance deep learning library," in *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. Red Hook, NY, USA: Curran Associates Inc., 2019.