Machine and Deep Learning for Coating Thickness Prediction using Lamb Waves

Maximilian Schmitz[a], Jin-Yeon Kim[a], Laurence J. Jacobs[a,b,*]

[a] *School of Civil and Environmental Engineering, Georgia Institute of Technology, Atlanta, GA, USA*
[b] *George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA, USA*

Abstract

This research applies machine and deep learning to nondestructively characterize the thickness and uniformity of a coating in a layered system using dispersion curves. A finite element analysis model is first used to computationally model transient, guided Lamb waves propagating in coated specimens with different coating thicknesses. These time-domain signals are then processed with a two-dimensional Fourier transform to obtain the corresponding frequency-wavenumber relation, which are the dispersion maps of the coated specimen. Dispersion maps are characteristic and depend on both the coating thickness and uniformity, plus its elastic properties (which are taken to be constant). Computationally simulated dispersion maps for a variety of coating properties are obtained and then further processed to extract a feature representation for each dispersion curve. Those extracted features are fed into machine learning classifiers which allow a thickness classification. This machine learning procedure is shown to be effective in classifying the thickness of a uniform coating. However, if the coating thickness is nonuniform, deep learning, specifically a convolutional neural network architecture, is used for classification. The network is evaluated and tested, and recommendations on its use are given.

Keywords: Machine Learning, Deep Learning, Inversion, Wave Propagation, Lamb Wave, Fourier Transform

1. Introduction

Thin coating layers or films on a substrate are often used to protect the substrate material from corrosion, wear, oxidation, melting, and thermal cycling. Examples include the thermal barrier coatings on gas turbine blades [1], vacuum-deposited thin films on microelectronic devices [2], and Chromium-coatings on zirconium alloy in nuclear fuel claddings [3]. Nondestructive evaluation (NDE) of thickness and mechanical properties of thin coating layers has been a richly investigated topic using different measurement techniques that are based on various physical principles, including magnetics, eddy current, X-ray photoelectron spectroscopy, and ultrasound. In many applications, achieving the coating thickness to be within a desired range is more important than the other properties to meet the intended function of the coating layer. Ultrasonic techniques have a few advantages among candidate NDE methods. For example, [4] employed normal incident ultrasound in the 10-20 MHz range to determine the thickness of a thin coating layer on a thick substrate. Their examples were for $50-100\mu$m thick epoxy and Plexiglas coatings on an aluminum substrate. [5] developed an acousto-optic technique based on the phase variation of incident and reflected ultrasound, while [6] analyzed this problem to establish the phase dependency on the coating parameters. [7] measured the thickness of a very thin (submicron) fluid lubrication layer between the outer raceway and ball in a ball bearing using high frequency (50 MHz) ultrasound. Other research used simulated or experimentally measured data and compared this data to a theoretical model [8], or used an inversion scheme based on the Global Matrix Method, as done by [9, 10]. These inversion approaches require a thorough understanding of the underlying wave propagation mechanics, including an accurate physics-based model.

In general, determining the thickness of thin coating layers is a challenging problem where the level of uncertainties about the thickness to be determined can be high. These uncertainties come from the complex physics of guided waves – multi modal and dispersive – so a systematic approach is needed to statistically describe the unknown coating thickness. This paper employs a data-driven, machine learning and deep learning model approach to solve the problem. The uncertainties are mitigated by determining and using a feature, the key physical characteristic for the thickness determination, which is strategically extracted from a map of the entire multi modal, dispersive wave physics of the problem.

By using this machine and deep learning model to determine the coating thickness of a coated specimen, only "sufficient" dispersion training data for a varying coating thickness of a given layered system is required. In this way, a data-driven, machine and deep learning model are used as surrogates for the physics-based model.

*Corresponding author
Email addresses: mschmitz7@gatech.edu (Maximilian Schmitz), jykim@gatech.edu (Jin-Yeon Kim), laurence.jacobs@coe.gatech.edu (Laurence J. Jacobs)

The layered system under consideration is a chromium coated, zirconium-4 plate. Note that the materials are assumed to be linear elastic with no damping. The training dispersion relations are computationally produced via finite element analysis (FEA) – FEA is used to simulate experimentally measured time-domain signals in physical specimens. Thus the forward problem is solved with an FEA model, specifically Abaqus/Explicit [11]. The effectiveness and validity of FEA for guided wave propagation is well defined and guidelines for mesh and time step lengths are considered [12, 13, 14].

The forward problem provides time-domain values for particle displacement. Dispersive and multi-modal Lamb wave signals in a coated plate specimen are complex, so digital signal processing (DSP) techniques are vital to extract dispersion information from these time-domain signals. Among others, candidate DSP methods include the Short-time Fourier transform [15], various wavelet transforms like the Gabor wavelet transform [16], and the two-dimensional Fourier transform (2D-FFT) [17]. This work relies on the 2D-FFT to calculate dispersion curves from these displacement signals which can be compared to analytically obtained dispersion curves.

A key contribution of this research is that only "sufficient" training examples from the forward problem are needed to solve the inverse problem. The idea is that these dispersion relations are sensitive to physical variations in coating thickness of a layered plate system. Depending on the coating thickness and material properties, either the coating layer or the base plate has a bigger influence on the dispersive behavior of the combined, layered system. An understanding of a system of a thin coating layer with similar properties to the base plate is analyzed, and specific features are manually identified to allow for inversion on the thin coating layer thickness.

If additional information on coating thickness uniformity is needed, the proposed machine learning algorithm approaches its limit. Then, a deep learning method can be used. Deep learning methods for wave inversion have been used in geophysics and geoscience, but to the authors' knowledge, have not yet been applied in a similar context to nondestructive evaluation. [18] describes different use cases for machine learning in seismology, while [19] developed a deep network consisting of encoder and decoder to obtain subsurface velocity structures for subsurface characterization in geoscience. Integrating known physical relationships into the training process has been conducted by [20, 21].

2. General Approach

This research investigates coating thickness and uniformity with guided Lamb waves. A summary of the solution approach to this problem is:

1. use FEA to simulate realistic experimental, time-domain ultrasonic signals;
2. take these time-domain signals to create dispersion curves (data) with a 2D-FFT and post-processing (clip values between 0 and 0.0001 and remove values smaller than median);
3. extract features through non-maximum suppression fit;
4. plug features into machine learning classifiers to train and determine coating thickness; and
5. show how uniformity can be classified with a deep learning network.

Note that inversion of this specific system is difficult to perform because the coating layer is very thin in comparison to the base plate thickness (coating thickness on the order of a hundredth of the plate's thickness) and that the elastic material properties of the coating and base plate components are comparably similar (coating's and plate's Young's modulus differ by a factor smaller than three). As a result, a novel contribution of this study is to show that data-driven coating thickness inversion is possible with machine learning combined with FEA simulations.

3. Specimen Description

Consider a model consisting of a zirconium alloy with a chromium coating proposed to improve the performance of accident tolerant fuels [22, 23], specifically a two-layered bonded system consisting of a 8 cm long and 1 mm thick zirconium-4 plate with a chromium coating of varying thickness between 10µm and 600µm. The mechanical properties of both materials are shown in Table 1. For the FEA simulation, the excitation source is located at the top left corner of the coated side of the coating-plate system, just as it would be for the case of a real ultrasonic measurement.

Table 1: Material properties for coating and plate

Material	Young's Modulus	Poisson's Ratio	Density
Zircaloy-4	$E_p = 99.3$GPa	$\nu_p = 0.37$	$\rho_p = 6.56 \frac{\text{g}}{\text{cm}^3}$
Chromium	$E_c = 279$GPa	$\nu_c = 0.21$	$\rho_c = 7.19 \frac{\text{g}}{\text{cm}^3}$

4. Analysis of Analytically Obtained Dispersion Curves

Dispersion curves are obtained analytically by solving the characteristic equations for a two-layer plate.[24] Figure 1 shows the dispersion curves for the coating, the plate, and the combined system. Since the purpose of this research is to obtain information about the coating layer, and not about the underlying zirconium base plate, the dominant physics of the coating layer needs to be extracted without the influence of the plate.

The idea is to be able to analyze the lower order modes of the coating without (or with little influence of) the

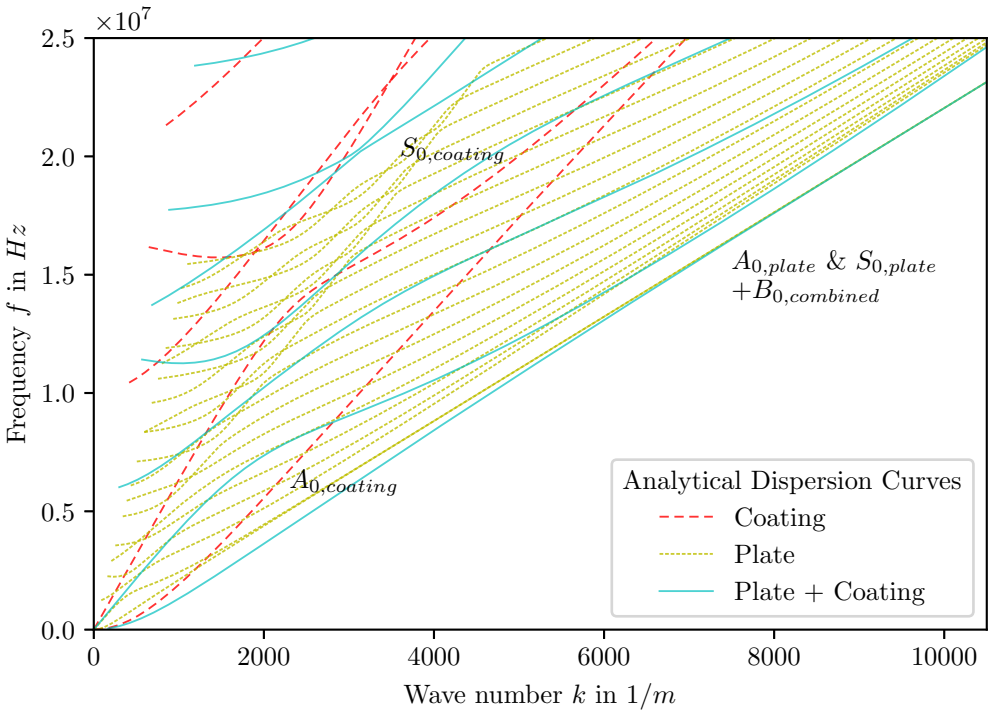

Figure 1: Analytical dispersion curves for a 200µm coating (red dashed), 1mm plate (green dashed), and the layered combined system (solid blue).

modes from the base plate. Since the coating thickness is very small, the idea is to analyze the change in dispersion curves for high frequencies, f with $f \in [10, 25]$ MHz and wavenumbers within the range of $k \in [500, 7000]\frac{1}{\text{m}}$. The reason for this is that for these frequencies and wavenumbers, the distance between the lower order modes of the coating (red), i.e. $A_{0,coating}$ and $S_{0,coating}$, and the lower order modes of the plate (green), i.e. $A_{0,plate}$ and $S_{0,plate}$, is maximized, while keeping the wavenumber and frequency reasonably low to meet requirements needed for practical ultrasonic measurements. Note that S_0 and A_0 designate, respectively the lowest symmetric and asymmetric modes of a specific plate/layer. It can be seen that in the proposed frequency and wavenumber range, the $S_{0,coating}$ and $A_{0,coating}$ are present, while only higher order modes of the plate show up. It should be noted that moving to higher frequencies and wavenumbers also means approaching the domain where Lamb waves have less influence, and Rayleigh waves, which concentrate their energy close to the boundary/surface, become more dominant.

Simulating the system using FEA is a second modelling option in addition to analytically solving the characteristic equations to obtain the dispersion curves of a given system. The behavior described above is seen in the FEA simulated dispersion maps like Fig. 2. Here, high-intensity values (blue to white) are registered close to the origin, because the lower order modes of the coating and of the base plate are close to each other, and because most of the energy is contained in low frequencies and wavenumbers. It can be seen that along the coating's mode $A_{0,coating}$ the intensity values are high for higher frequencies and wavenumbers, and that the modes of the base plate are obvious only near this area. This means that the location where the coating's A_0 mode – and also when the higher-order modes of the plate show up – is sensitive to the coating thickness. The area for wavenumbers below $k = 2000$ 1/m is not as sensitive to the coating's thickness and therefore, is not further analyzed in this work, even though it still might contain useful information.

5. Forward Problem (Simulation)

A triangular displacement impulse is applied at the upper left corner of the coating in the FEA model so that frequencies up to $f_{max} = 628$ MHz are excited which enables a high resolution in the frequency domain as shown in Fig. 2. All nodes at the left edge of the system are constrained in the x_1 (in-plane) direction. A sketch of the FEA used models can be found in Appendix A. Since the objective of this research is to obtain a dispersion frequency-wavenumber representation for multiple frequencies, a wide range of frequencies needs to be excited. This is done with a triangular pulse, since the Fourier transform of a triangular pulse is the squared Sinus cardinal (Sinc) function.

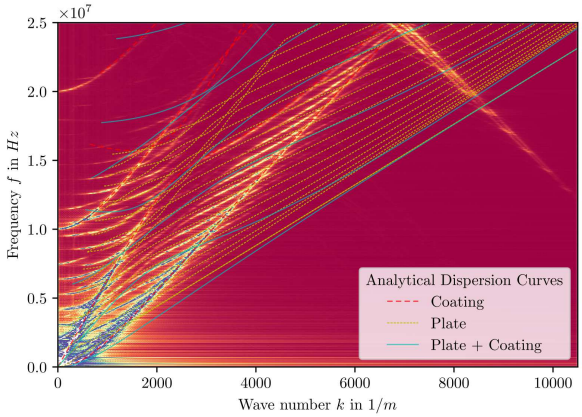

Figure 2: Analytical dispersion curves and simulated dispersion map for a coating of 200µm thickness. Blue values indicate a high intensity while red indicate a low intensity. The yellow line with negative gradient in the top right corner is connected to reflections and is not part of this discussion.

Sampling is conducted both in the time and spatial domains. The sampled displacement data over time and space are then transformed into the frequency-wavenumber domain using the 2D-FFT; the trajectories of high intensity ridges in this map correspond to the dispersion curves of the Lamb modes in the two layer plate. We call this 2D-representation a dispersion map.

When comparing the FEA simulated dispersion maps from different coating thicknesses (and constant base plate thickness) there are three main entities to look at: the gradient of the dominant lower order coating modes; the convergence/divergence of the lower-order coating modes; and the appearance of lower- and higher-order modes of the plate.

Figure 3 shows dispersion maps of three different coating thicknesses. The gradient of the dominant lower order coating modes, i.e. $S_{0,coating}$ and $A_{0,coating}$ which are shown in Fig. 3a, is increasing with thickness, while both modes converge in Fig. 3b. Additionally, with increasing thickness of the coating, lower- and higher-order modes of the plates disappear, such that almost no plate modes can be seen in Fig. 3c.

6. Inverse Problem via Machine Learning

Using this specific guided Lamb wave knowledge, the simulated dispersion curves can now be characterized and used for a machine learning-based inversion procedure. Here dispersion curve features are manually selected, with the goal of determining if a previously unseen dispersion curve belongs to a layered system whose coating layer satisfies the conditions of *thick enough* or *not thick enough*.

6.1. Feature Extraction

The feature extraction consists of two steps: non-maximum suppression to extract the coordinates of high-intensity values around the modes; and fitting a linear function to these points. The 2D-FFT serves as a two-dimensional map of intensity values in the frequency-wavenumber domain for a given specimen. The Lamb wave modes are encoded in the 2D-FFT spectrum as coherent points of high intensity. To extract these points, non-maximum suppression (NMS) is used [25]. NMS is a method from computer science and computer vision that presents a solution to extracting the coordinates of these local intensity maxima.

The extracted maximum intensity values from the dispersion curves from the previous section can now be used as an input to a linear fitting algorithm [26]

$$f(k) = a\,k + b, \tag{1}$$

with the gradient a and the y-intercept b. An example of a fit is shown in Fig. 4.

This approach works well for the coating thickness range between 10µm and 600µm, and this approach is not tested for thicknesses outside this range. The gradient is especially sensitive to a change in the coating's thickness. The physical explanation for this is:

1. the intensity is higher because the $A_{0,coating}$ mode is asymmetric, the coating layer is geometrically asymmetric, and the excitation is asymmetric;
2. the slope is the parameter directly related to the $A_{0,coating}$ mode velocity; and
3. for thinner coatings, $A_{0,coating}$ is well isolated in the dispersion map, and is not contaminated by the other modes. For thicker coatings, $A_{0,coating}$ is merges together with $S_{0,coating}$.

6.2. Applying Machine Learning Classifiers

A coated specimen is classified as, i.e. the labels are, *thick enough* or *not thick enough*, respectively. Exemplary and without loss of generality, a coating thickness of $h_{coating} =$ 200µm is chosen to be the threshold between *thick enough* and *not thick enough*. The scheme is tested with other thresholds and works analogously.

The extracted gradient and y-intercept for each dispersion map are now used as features for thickness classification. Since only two features are used, a planar visualization of the complete feature space is possible in Fig. 5. Each data point corresponds to one coated specimen with respective coating thickness and is color-coded according to its ground truth thickness classification.

The first step of applying a machine learning classifier is to normalize the input data. The feature values are normalized to the range between zero and one. The data obtained can now be fed into various classifiers. The classifiers used are k-Nearest Neighbor (kNN) [27], single layer perceptron [28], Support vector machines (SVM) with and without radial basis function kernel [29, 30], Gaussian processes (GP) [31] and feedforward networks (MLP) [32, 33]. The algorithms are then evaluated with k-fold cross-validation (CV) for $k = 5$ [34]. The results after hyperparameter tuning are shown in Table 2.

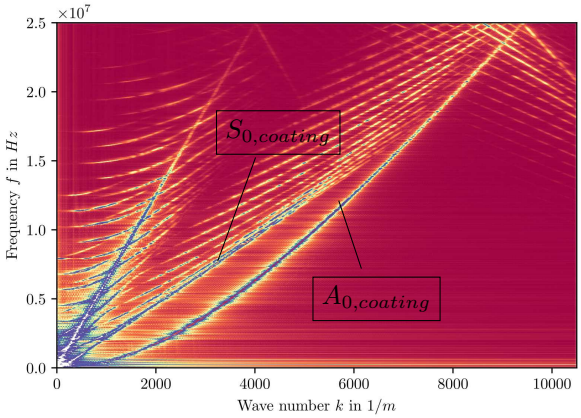

(a) Coating thickness of $h_{coating} = 40\mu m$

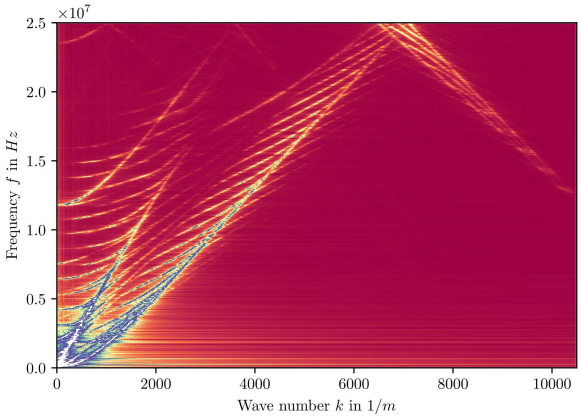

(b) Coating thickness of $h_{coating} = 170\mu m$

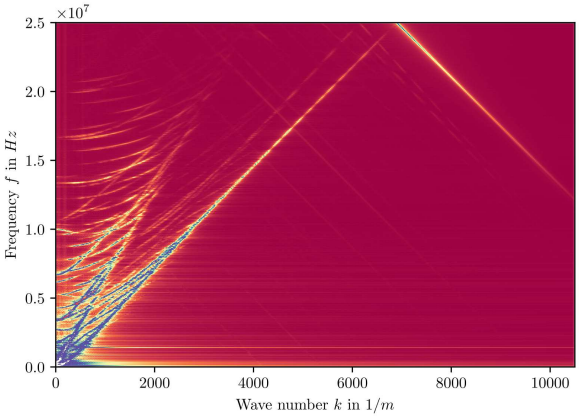

(c) Coating thickness of $h_{coating} = 600\mu m$

Figure 3: FEA simulated dispersion maps for selected thicknesses

Table 2: Comparison of accuracy and standard deviation after 5-fold cross-validation and hyperparameter tuning.

Classifier	Accuracy score	σ after CV
1-Nearest Neighbor	0.903	0.123
5-Nearest Neighbor	0.917	0.107
Perceptron	0.94	0.074
Linear Support Vector Machines	0.947	0.066
Radial Basis Function SVM	0.954	0.062
Gaussian Processes	0.903	0.132
Feedforward Network	0.917	0.107

It can be seen that all classifiers perform well given the dataset of 134 simulated uniformly coated plates, and achieve high accuracy of over 90%. Support Vector machines, especially support vector machines with a radial basis function kernel (RBFSVM) tend to have the highest accuracy with the smallest standard deviation after cross-validation, while one-nearest neighbor (1NN) and Gaussian processes do not perform as well.

This provides a visualization of two selected classifiers with the lowest (1NN in Fig. 6) and highest (RBFSVM in Fig. 7) accuracy, where the red area describes where a feature combination would be labeled as *not thick enough*, and the green area is where a feature combination would be classified as *thick enough*.

7. Inverse Problem via Deep Learning

The machine learning approach for inversion works well for a uniform thickness, but reaches its limits when a non-uniformity is incorporated in the coating of a simulated specimen. Comparing the dispersion curves of a non-uniform coating with its uniform counterpart, for an intended coating thickness of $h_{coating} = 200\mu m$, a depth of the gap as the non-uniformity of $h_{gap} = 100\mu m$ (which is a reduction of the coating thickness by half) over a length of $l_{gap} = 1.2$ cm, While there is a difference in the dispersion curves, neither a fitted gradient or y-intercept can be used to characterize this change. This same trend holds for other coating thicknesses and gap depths, indicating that the proposed machine learning-based inversion process is not sensitive to a non-uniformity in the coating.

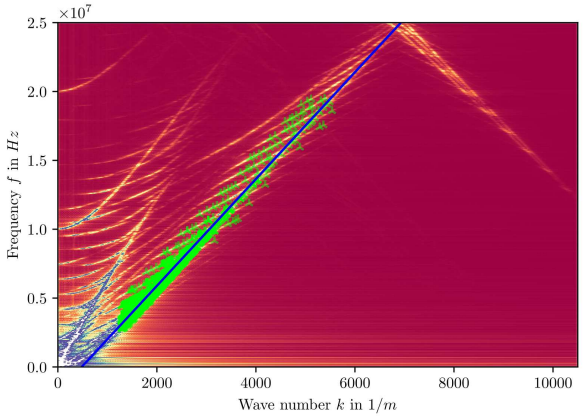

Figure 4: Coating thickness of $h_{coating} = 200\mu m$ with NMS suppression and fitted function.

This machine learning approach fails in this application mainly because any changes in dispersion curves due to non-uniformity are minor and do not follow a recognizable pattern. Instead, a convolutional neural network (CNN) approach is proposed to learn the non-uniformity features.

Feature classification learning by the CNN alone requires significantly more training data than the simple machine learning classifiers. To obtain the data for CNN training, 414 simulations with varying parameters were conducted on the Georgia Tech PACE cluster [35]. An overview of the simulation space can be found in Fig. 8. Each dot in this figure represents one simulation of a specimen with a given coating thickness $h_{coating}$ and depth of the gap h_{gap} as a measure of the extent coating layer non-uniformity. The gap length is not shown in the figure and varies between 14mm and 18mm. Simulations with a non-uniformity, i.e. a gap depth bigger than zero, are coded with red circles while uniform simulations are coded with a green diamond.

The first step in the development of a neural network is specifying the shape of the data at the input side. The dispersion map is first stored in the PNG format. This is feasible since PNG supports lossless compression [36]. With this procedure, the size of the input data can be reduced by a factor of 30. Since PNG files are generally referred to as *images*, this work will use the word *images* to refer to the dispersion curve inputs to the neural network too.

The PNG images created are then randomly shuffled and assigned to the training and the evaluation sets, such that 70% of the data is used for training and the remaining 30% are used for testing. Since the amount of data available, i.e. the number of simulations conducted, is fairly limited, the evaluation set consists of the test set. As with most neural networks, the proposed network in this work uses batches of input images for training – this study uses a batch size of 32. The process of loading a random batch of images, passing them through the network, calculating the loss of each image, differentiating loss with respect to the network weights, and updating them accordingly, is referred to as one *epoch* of training.

After conversion to a one-channel image, the image is resized to a rectangular input of a given resolution. This research proposes an input resolution of 1024×1024 to capture even small changes in the image, but smaller input resolutions might be feasible depending on the given problem. Note that the input resolution is a critical hyperparameter of the learning process. Using a resolution that is too small removes too much high-frequency information in the image, while too big a resolution slows down the learning process.

7.1. *SimpleWaveInvNet*

SimpleWaveInvNet is a comparable simple network for thickness inversion developed in this research to capture small changes in the input by using the smallest number of parameters possible [37]. It consists of three layer parts:

- Convolutional layers, consisting of four 2D convolutional layers, 2D max-pooling after the first, third, and fourth convolutional layer, ReLU[1] as activation function after each convolutional layer, a dropout layer in front of the last convolutional layer and a 2D batch norm after the last convolutional layer.
- Average pooling layer which conducts 2D adaptive average pooling to a 5×5 output. This layer allows for the use of variation in the input size to the network without the need of restructuring the entire network architecture.
- Fully connected layers, comprising three fully connected layers with a dropout layer in front of the first and second fully connected layer, and a ReLU layer behind these two layers.

The loss function for the SimpleWaveInvNet is chosen to be the negative log-likelihood, or *NLLLoss* [40].

SimpleWaveInvNet is trained over 200 epochs. Fig. 9 shows the training loss history (a) and training accuracy history (b) for both the training (blue) and the validation set (red). It can be seen that the training error for both the training and validation sets deceases quickly in the beginning and then decreases more slowly until the end of the training process. The shape of the loss curves complies with the expected learning curves of a neural network.

The accuracy is increasing heavily in the beginning and then reaches a plateau, for both the training and validation sets. Again, the overall learning curve complies with the expected shape. Note that the validation accuracy varies more than the accuracy of the training set, which accompanies the high noise in the validation loss. A reason for

[1] ReLU stands for rectified linear unit and is often used as a default activation function in deep learning [38, 39]. The pointwise nonlinear function f is defined as $f(x) = \max(0, x)$.

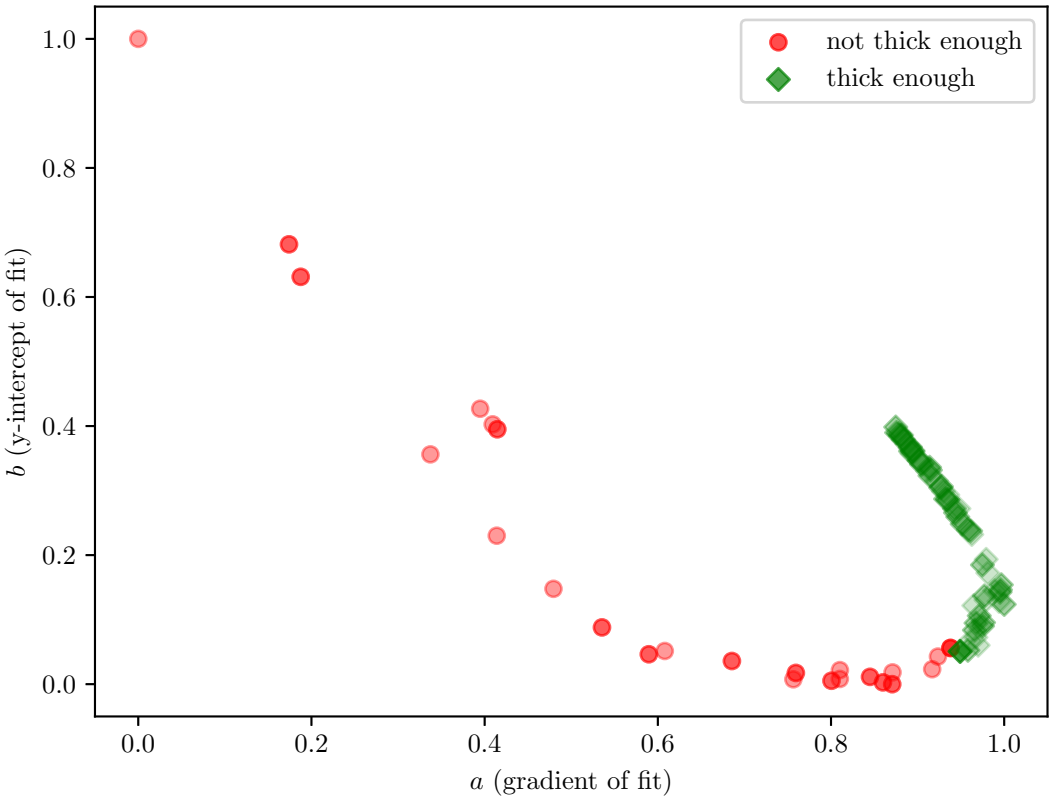

Figure 5: Feature plot for dispersion inversion with fitted function of shape $f(k) = a\,k+b$ and labeling according to threshold $h_{coating} = 200\mu m$.

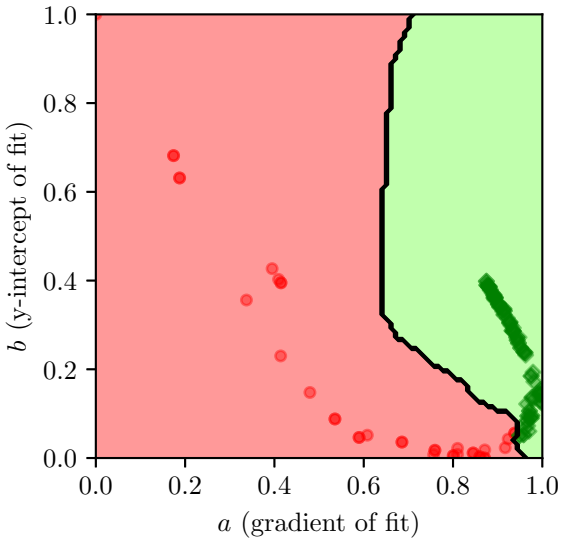

Figure 6: kNN classifier with $k = 1$.

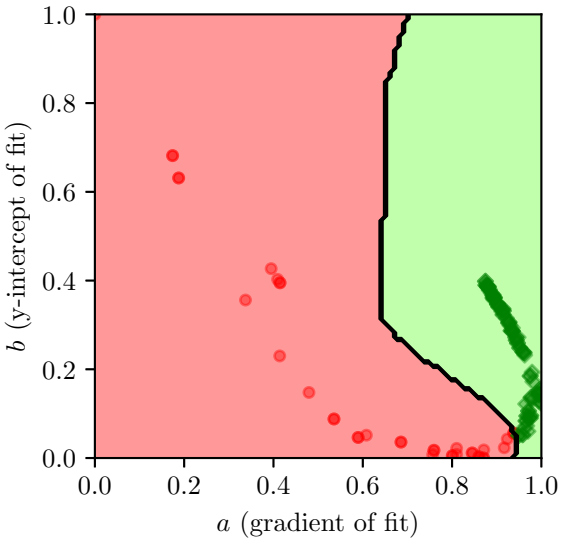

Figure 7: SVC classifier with RBF kernel and regularization parameter `C` = 4 and kernel coefficient `gamma` = 2.

this might be the limited data set size as well as a batch size chosen to be too small. Some validation batches might not be a good representation of the overall dataset.

Figure 10 is a confusion matrix from the evaluation set for SimpleWaveInvNet. This confusion matrix shows the ground truth labels on the vertical and the predicted labels on the horizontal axis. For SimpleWaveInvNet, the majority of the non-uniform samples are classified correctly. Classifying samples with a uniform coating is not as unambiguous as for the non-uniform coating, but it still classifies 76% correctly. This demonstrates that SimpleWaveInvNet is able to provide a uniform/non-uniform thickness classification.

Figure 11 is a visualization of how SimpleWaveInvNet performs on a test set. Most of the data are classified correctly, while most false positive classifications are above a thickness of 300µm. This is exactly the third set of thicknesses described in Section 5. At this coating layer thickness, a non-uniformity in the coating has a smaller impact. It is expected that even with more training data, the classification of a uniformity is approaching a limit with the approach given in this research.

In addition to SimpleWaveInvNet, ResNet18 [41] is also tested on the same problem, ResNet18 performed poorly since the number of training data is not sufficient for the large number of parameters needed for this network.

8. Conclusion

This research applies machine and deep learning to nondestructively characterize the thickness and uniformity of a coating in a layered system using dispersion curves. FEA is used for the forward problem to computationally model transient, guided Lamb waves propagating in coated specimens with different coating thicknesses. These time-domain signals are then processed with a 2D-FFT to obtain the corresponding dispersion maps of the coated specimens and then further processed to extract a feature representation for each dispersion curve. The inversion procedure to determine the coating layer thickness is accomplished by feeding these extracted features into machine learning classifiers. This machine learning procedure is shown to be effective in classifying the thickness of a uniform coating.

However, if the coating thickness is nonuniform, deep learning, specifically a CNN network architecture, is used for classification. This deep learning-based approach demonstrated that a CNN is capable of learning if a given specimen contains a non-uniformity in the coating or not.

Recommendations for future work include increasing the training dataset for a better representation to input into the CNN. Since data besides the $A0$ and $S0$ do not contain physically meaningful information, learning an additional neural network as an encoder could decrease the input size of the CNN and increase training performance. Finally, current research is examining experimental ultrasonic measurements on a physical layered plate system.

9. Acknowledgement

This research is partially supported by the Nuclear Energy University Program (NEUP) under project 20-19660, through research cyberinfrastructure resources and services provided by the Partnership for an Advanced Computing Environment (PACE) at the Georgia Institute of Technology, Atlanta, Georgia, USA, and by the Baden-Württemberg Foundation.

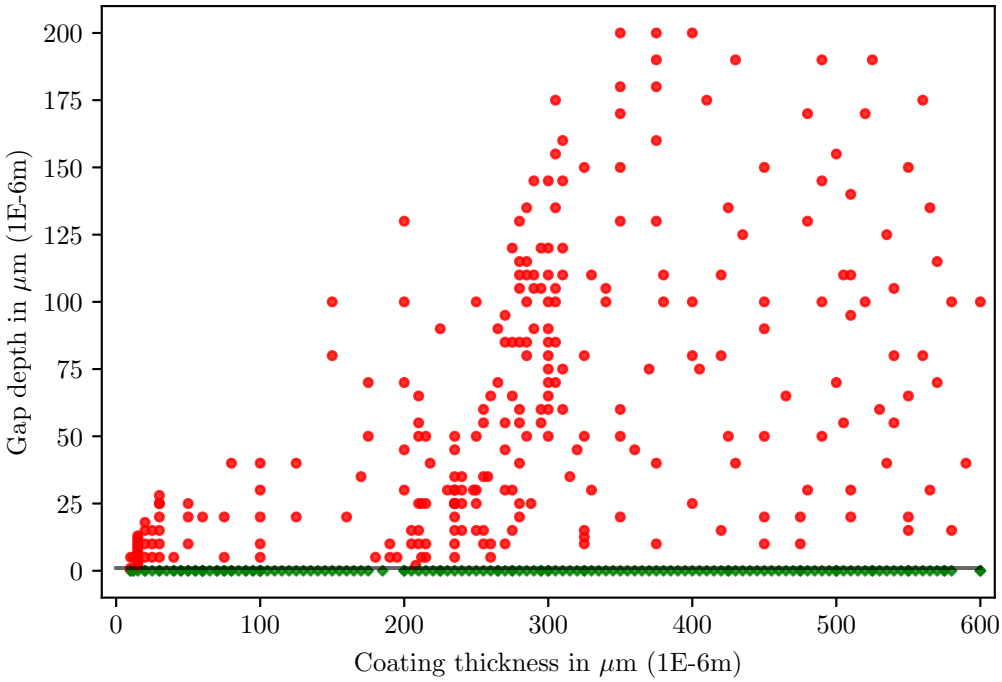

Figure 8: Simulation space for randomly selected 414 simulations with 148 (= 35.75%) simulations which are uniform and 266 (= 64.25%) simulations with non-uniformity.

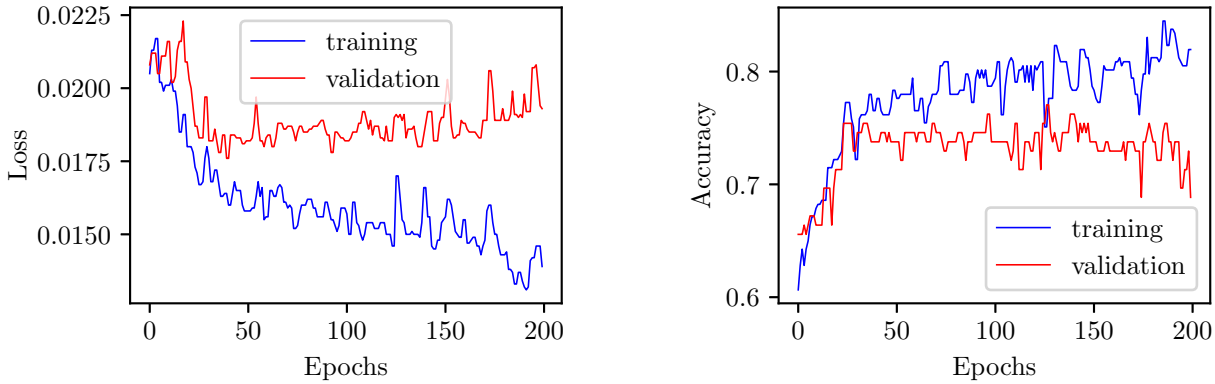

(a) Training loss history.

(b) Training accuracy history.

Figure 9: SimpleWaveInvNet training performance after 200 epochs and median filtering.

References

[1] J. H. Perepezko, The Hotter the Engine, the Better, Science 326 (2009) 1068–1069.

[2] Z. Bao, A. J. Lovinger, A. Dodabalapur, Highly ordered vacuum-deposited thin films of metallophthalocyanines and their applications in field-effect transistors, Advanced Materials 9 (1997) 42–44.

[3] S. J. Zinkle, K. A. Terrani, J. C. Gehin, L. J. Ott, L. L. Snead, Accident tolerant fuels for LWRs: A perspective, Journal of Nuclear Materials 448 (2014) 374–379.

[4] V. K. Kinra, C. Zhu, Ultrasonic nondestructive evaluation of thin (sub-wavelength) coatings, The Journal of the Acoustical Society of America 93 (1993) 2454–2467.

[5] S. Devolder, M. Wevers, P. De Meester, O. Leroy, Thin layer thickness measurements based on the acousto-optic technique, Applied Physics Letters 68 (1996) 1732–1734. Publisher: American Institute of Physics.

[6] F. W. Windels, S. Vanaverbeke, O. Leroy, Thin coating characterization by Rayleigh waves: An analytical model based on normal-mode theory, The Journal of the Acoustical Society of America 110 (2001) 1349–1359. Publisher: Acoustical Society of America.

[7] J. Zhang, B. W. Drinkwater, R. S. Dwyer-Joyce, Acoustic measurement of lubricant-film thickness distribution in ball bearings, The Journal of the Acoustical Society of America 119 (2006) 863–871. Publisher: Acoustical Society of America.

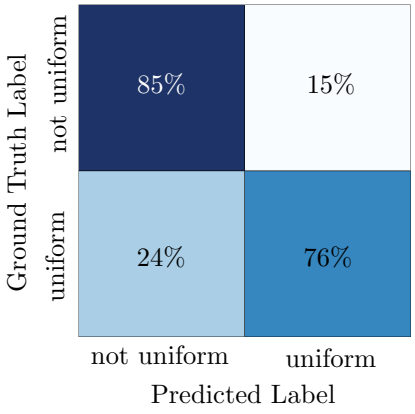

Figure 10: Confusion matrix for SimpleWaveInvNet.

[8] J. Stolzenburg, J. W. Doane, J. Jarzynski, L. J. Jacobs, Near field inversion method to measure the material properties of a layer, NDT & E International 36 (2003) 523–533.

[9] J. Koreck, C. Valle, J. Qu, L. J. Jacobs, Computational Characterization of Adhesive Layer Properties Using Guided Waves in Bonded Plates, Journal of Nondestructive Evaluation 26 (2007) 97–105.

[10] J. Koreck, Computational characterization of adhesive bond-properties using guided waves in bonded plates, Master's thesis, School of Civil and Environmental Engineering, Georgia Institute of Technology, 2006.

[11] M. Smith, ABAQUS/Standard User's Manual, Version 6.14, Dassault Systéms Simulia Corp, United States, 2014.

[12] F. Moser, L. J. Jacobs, J. Qu, Modeling elastic wave propagation in waveguides with the finite element method, NDT & E International 32 (1999) 225–234.

[13] R. Seifried, L. J. Jacobs, J. Qu, Propagation of guided waves in adhesive bonded components, NDT & E International 35 (2002) 317–328.

[14] C. Valle, M. Niethammer, J. Qu, L. J. Jacobs, Crack characterization using guided circumferential waves, The Journal of the Acoustical Society of America 110 (2001) 1282–1290.

[15] M. Niethammer, L. J. Jacobs, J. Qu, J. Jarzynski, Time-frequency representations of Lamb waves, The Journal of the Acoustical Society of America 109 (2001) 1841–1847.

[16] F. Lanza di Scalea, P. Rizzo, A. Marzani, Propagation of ultrasonic guided waves in lap-shear adhesive joints: Case of incident a0 Lamb wave, The Journal of the Acoustical Society of America 115 (2004) 146–156.

[17] D. Alleyne, P. Cawley, A two-dimensional fourier transform method for the measurement of propagating multimode signals, The Journal of the Acoustical Society of America 89 (1991) 1159–1168.

[18] Q. Kong, D. T. Trugman, Z. E. Ross, M. J. Bianco, B. J. Meade, P. Gerstoft, Machine learning in seismology: Turning data into insights, Seismological Research Letters 90 (2018) 3–14.

[19] Y. Wu, Y. Lin, Inversionnet: An efficient and accurate data-driven full waveform inversion, IEEE Transactions on Computational Imaging 6 (2020) 419–433.

[20] R. Rojas-Gómez, J. Yang, Y. Lin, J. Theiler, B. Wohlberg, Physics-consistent data-driven waveform inversion with adaptive data augmentation, IEEE Geoscience and Remote Sensing Letters 19 (2020) 1–5.

[21] Y. Ren, X. Xu, S. Yang, L. Nie, Y. Chen, A physics-based neural-network way to perform seismic full waveform inversion, IEEE Access 8 (2020) 112266–112277.

[22] C. Tang, M. Stueber, H. J. Seifert, M. Steinbrueck, Protective coatings on zirconium-based alloys as accident-tolerant fuel (ATF) claddings, Corrosion Reviews 35 (2017) 141–165. Publisher: De Gruyter.

[23] D. V. Nguyen, M. Le Saux, L. Gélébart, J.-C. Brachet, J.-P. Bonthonneau, A. Courcelle, R. Guillou, E. Rouesne, S. Urvoy, Mechanical behavior of a chromium coating on a zirconium alloy substrate at room temperature, Journal of Nuclear Materials 558 (2022) 153332.

[24] A. Huber, Dispersion calculator, version 1.11, 2021. German Aerospace Center (DLR).

[25] A. Neubeck, L. Van Gool, Efficient Non-Maximum Suppression, in: 18th International Conference on Pattern Recognition (ICPR'06), volume 3, 2006, pp. 850–855. doi:10.1109/ICPR.2006.479.

[26] G. Chavent, Nonlinear Least Squares for Inverse Problems: Theoretical Foundations and Step-by-Step Guide for Applications, Scientific Computation, Springer Netherlands, Dordrecht, 2010. URL: http://link.springer.com/10.1007/978-90-481-2785-6. doi:10.1007/978-90-481-2785-6.

[27] K. Taunk, S. De, S. Verma, A. Swetapadma, A Brief Review of Nearest Neighbor Algorithm for Learning and Classification, in: 2019 International Conference on Intelligent Computing and Control Systems (ICCS), 2019, pp. 1255–1260. doi:10.1109/ICCS45141.2019.9065747.

[28] F. Rosenblatt, The perceptron: A probabilistic model for information storage and organization in the brain., Psychological Review 65 (1958) 386–408.

[29] A. Patle, D. S. Chouhan, SVM kernel functions for classification, in: 2013 International Conference on Advances in Technology and Engineering (ICATE), 2013, pp. 1–9. doi:10.1109/ICAdTE.2013.6524743.

[30] B. E. Boser, I. M. Guyon, V. N. Vapnik, A training algorithm for optimal margin classifiers, in: Proceedings of the fifth annual workshop on Computational learning theory, COLT '92, Association for Computing Machinery, New York, NY, USA, 1992, pp. 144–152. URL: https://doi.org/10.1145/130385.130401. doi:10.1145/130385.130401.

[31] C. E. Rasmussen, Gaussian Processes in Machine Learning, in: O. Bousquet, U. von Luxburg, G. Rätsch (Eds.), Advanced Lectures on Machine Learning: ML Summer Schools 2003, Canberra, Australia, February 2 - 14, 2003, Tübingen, Germany, August 4 - 16, 2003, Revised Lectures, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 2004, pp. 63–71. URL: https://doi.org/10.1007/978-3-540-28650-9_4. doi:10.1007/978-3-540-28650-9_4.

[32] I. Goodfellow, Y. Bengio, A. Courville, Deep learning, The MIT Press, Cambridge, Massachusetts, 2016.

[33] K. P. Murphy, Machine Learning: A Probabilistic Perspective, The MIT Press, 2012.

[34] P. Refaeilzadeh, L. Tang, H. Liu, Cross-Validation, in: L. LIU, M. T. ÖZSU (Eds.), Encyclopedia of Database Systems, Springer US, Boston, MA, 2009, pp. 532–538. URL: https://doi.org/10.1007/978-0-387-39940-9_565. doi:10.1007/978-0-387-39940-9_565.

[35] PACE, Partnership for an Advanced Computing Environment (PACE), Georgia Institute of Technologyy, 2017. URL: http://www.pace.gatech.edu.

[36] D. Salomon, Data Compression: The Complete Reference, 4th ed ed., Springer, London, 2007.

[37] M. Schmitz, Deep Learning in Ultrasonic Wave Inversion for Thin Coatings, MS thesis, Georgia Institute of Technology, Atlanta, GA, 2022.

[38] V. Nair, G. E. Hinton, Rectified linear units improve restricted boltzmann machines, in: ICML, 2010, pp. 807–814. URL: https://icml.cc/Conferences/2010/papers/432.pdf.

[39] M. Zeiler, M. Ranzato, R. Monga, M. Mao, K. Yang, Q. Le, P. Nguyen, A. Senior, V. Vanhoucke, J. Dean, G. Hinton, On rectified linear units for speech processing, in: 2013 IEEE International Conference on Acoustics, Speech and Signal Processing, 2013, pp. 3517–3521. doi:10.1109/ICASSP.2013.6638312.

[40] D. Zhu, H. Yao, B. Jiang, P. Yu, Negative Log Likelihood Ratio Loss for Deep Neural Network Classification, in: Proceedings of the Future Technologies Conference (FTC) 2019,

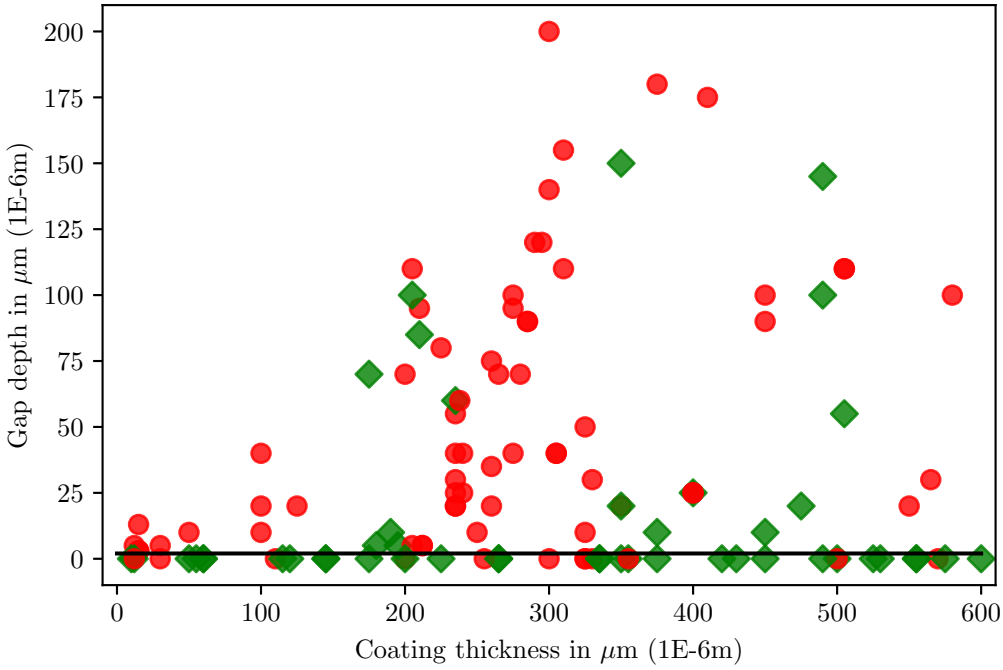

Figure 11: Classification of SimpleWaveInvNet for test set of 122 simulations with labels provided by network (green diamonds = uniform) and (red circles = not uniform).

Springer International Publishing, 2019, pp. 276–282. URL: http://arxiv.org/abs/1804.10690, arXiv: 1804.10690.

[41] K. He, X. Zhang, S. Ren, J. Sun, Deep Residual Learning for Image Recognition, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778. URL: http://arxiv.org/abs/1512.03385. doi:10.1109/CVPR.2016.90.

Appendix A. Finite Element Model

The simulation of the FEA model is conducted in Abaqus [11]. The triangular excitation is applied on the top left node of the system, which reaches its maximum extent after t_{max} = 20ns. A sketch of the simulation model with uniform coating is shown in Fig. A.12a, and a sketch of the model with non-uniform coating in Fig. A.12b.

Sampling in the time and the spatial domains is conducted every ΔT = 20ns and every ΔX = 6µm, respectively. Simulating one of the models on the Georgia Tech high-performance cluster PACE [35] takes between 15 and 25 hours depending on the thickness and geometry of the coating.

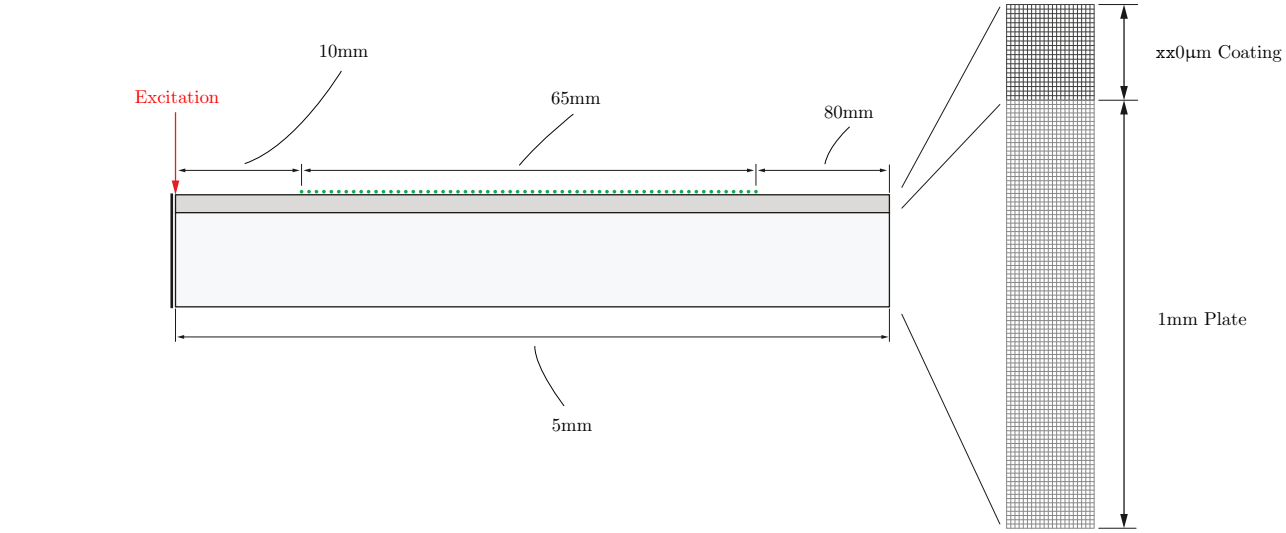

(a) FEA model with uniform coating (not drawn to scale).

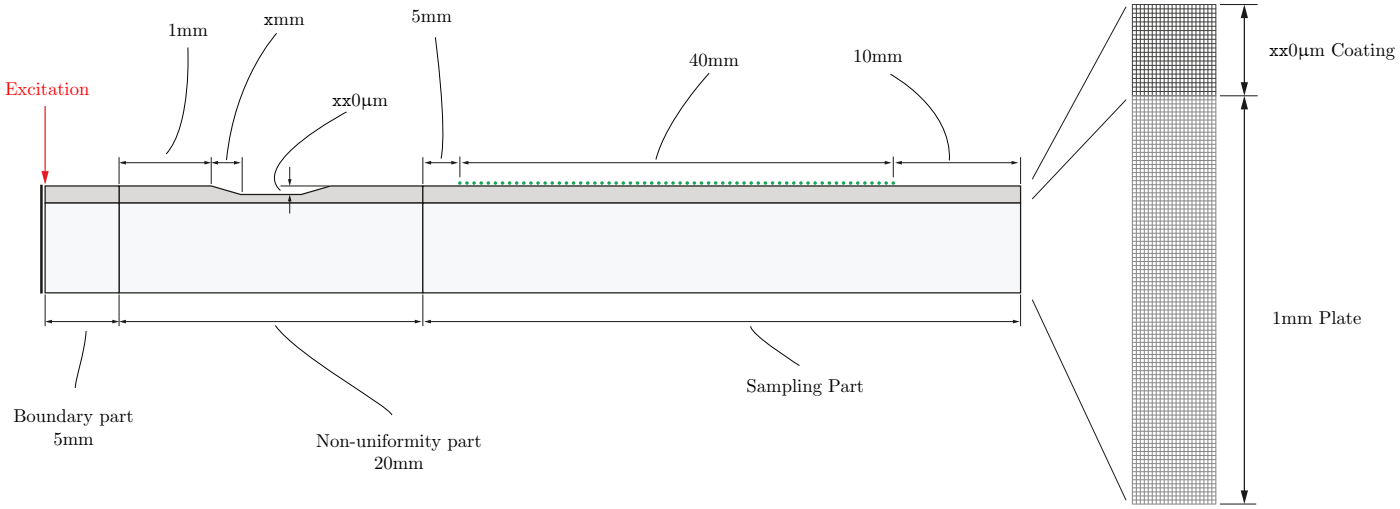

(b) FEA model with non-uniform coating (not drawn to scale).

Figure A.12: Sketches of both simulation models. Elements with an `x` specify parameters which are varied.