



Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

ENHANCING ACOUSTIC METAMATERIAL PERFORMANCE AND FABRICATION ROBUSTNESS WITH DEEP LEARNING

OLUWASEYI OGUN

SUPERVISOR: DR JOHN KENNEDY

Department of Mechanical, Manufacturing & Biomedical Engineering

School of Engineering

Trinity College Dublin, *the University of Dublin*

D02 PN40

Ireland

September, 2026

A dissertation submitted to the University of Dublin in partial
fulfilment of the requirements for the degree of PhD

Declaration

I agree that this thesis was completed in line with the plagiarism provisions in the General Regulations of the University Calendar for the current year, found at <http://www.tcd.ie/calendar>.

I have completed the Online Tutorial on avoiding plagiarism ‘Ready Steady Write’, located at <http://tcd-ie.libguides.com/plagiarism/ready-steady-write>.

I agree that this thesis will not be publicly available but will be available to TCD staff and students in the University’s open access institutional repository on the Trinity domain only.

Oluwaseyi Ogun,
February 2026

Abstract

Conventional acoustic materials, such as porous or fibrous absorbers, are limited in their ability to achieve low-frequency sound attenuation within compact volumes. Acoustic meta-materials (AMMs) offer a promising alternative, enabling extreme subwavelength resonances and tunable acoustic performance. However, designing AMMs that are simultaneously high-performing, fabrication-robust, and physically interpretable remains a major challenge. Additive manufacturing (AM) allows complex geometries to be realised, but introduces imperfections that can degrade acoustic performance, while traditional design methods often provide limited insight into the underlying structure–performance relationships.

In this thesis, robustness is defined as the ability of an AMM design or predictive model to maintain consistent acoustic performance in the presence of manufacturing tolerances, modeling uncertainties, and experimental variability. Interpretability refers to the capacity to extract physically meaningful insights from data-driven models, enabling designers to understand how geometric features influence acoustic response rather than treating models as black boxes.

This thesis presents a comprehensive framework for the design, optimisation, and robust realisation of AMM absorbers using deep learning (DL) approaches, with an emphasis on interpretability and real-world manufacturability. An AMM absorber, termed the Hybrid Resonant Acoustic Meta-Absorber (HRAM), consisting of an Archimedean spiral embedded in a hexagonal cavity with a perforated top plate, was developed and successfully validated through analytical modeling, high-fidelity numerical simulations, and experimental testing. The analytical model, while sufficient for the subsequent data-driven studies, showed observable discrepancies when compared to experimental results especially in terms of bandwidth characteristic. The high-fidelity numerical simulation also shows observable bandwidth mismatches relative to experimental curves. These observable deviations were largely attributed to unmodeled physical effects in the absorber behaviour, manufacturing issues and the un-

realistic assumption of sound-hard boundaries in numerical simulations. Furthermore, due to the computational demand of the full-wave (FW) numerical simulation, a simplified numerical model (SNM) was developed which significantly reduced the simulation time from the order of hours to minutes, and this simplified model was subsequently used in all references to numerical results.

Using the validated analytical model, a representative spectrum–geometry dataset was generated and employed to compare traditional optimisation methods, exemplified by a generalised pattern search (GPS), with deep learning (DL) approaches. A structured latent autoencoder (SLAE) was trained for both forward and inverse design tasks, and its performance was enhanced using on-the-fly data augmentation (artificially increasing the size and diversity of the training dataset by applying spectral transformations to existing data) and a transfer learning (TL) strategy incorporating sparse experimental data to capture manufacturing variability. The resulting SLAE-TL model demonstrated efficient, robust, and predictive performance, bridging the gap between theoretical design and real-world behaviour.

Beyond predictive capability, a novel design-space segmentation approach was developed to provide interpretability, revealing the relationship between geometric features and acoustic performance. This framework was applied to the design of a broadband AMM absorber for room acoustics, demonstrating both practical utility and scientific insight.

Collectively, the work establishes a unified methodology for the design of robust, interpretable, and high-performance AMMs. By integrating physics-based modelling, deep learning, and experimental validation, this research addresses long-standing challenges in AMM design and paves the way for real-world deployable solutions, advancing both the scientific understanding and practical implementation of AMMs.

To potent mix of curiosity, “what if?”, and “this might just work!”, for fueling research ideas. To procrastination, for giving me the time to explore every possible distraction before finally sitting down to write.

To my laptop, for surviving countless near-crashes. To Ctrl+Z, for undoing every regrettable mistake, and Ctrl+S, for saving me from the abyss of despair more times than I can count. To my family and friends, who believed in me when I thought I couldn’t. To my supervisor, who patiently endured my overuse of metaphors, complex charts, and moments of existential crisis. And ultimately to God, for the hope of a Phd: In thy hands, o Lord!

This thesis exists because of you all, and because deadlines are a surprisingly effective motivator.

Acknowledgements

They say it takes a village to raise a child. I have discovered it also takes a small city, a well-equipped laboratory, and a suspicious amount of caffeine to raise a PhD. I have been fortunate to have an extraordinary support team making sure I didn't drown, burn myself, or run in the wrong direction entirely.

First and foremost, I am deeply grateful to my supervisor, Dr. John Kennedy, for his good humour, unwavering support, and patient guidance throughout this entire adventure. His ability to combine high-level insight with a calm “don't worry, it's fixable” attitude kept both my research and my sanity intact. Thank you for giving me so many unique opportunities, for encouraging independence while offering a safety net, and for being, in all sincerity, the best supervisor I could have hoped for.

I extend sincere thanks to Prof. Henry Rice for his generous willingness to share knowledge, always delivered with clarity, patience, and the occasional gentle nudge in the right direction. Your advice has been invaluable.

To the technical staff at Trinity College Dublin — Mick, Brendan, and Alex — thank you for your expertise with the 3D printing facilities, and for ever refraining to say, “another batch of prints with no known expiration date” when I showed up with strange, overly ambitious CAD files.

Special thanks to my wonderful colleagues in the fluids lab, Agnieszka Ciochon, Karina Einicke, Eoghan Ross, Morgane Bellec, Lara Flanagan, Jiayu Wang, and Md Robiul Hossain, for your camaraderie, the countless get-togethers, the laughter, and those legendary games (which I'm sure future historians will struggle to categorise). A particular shout-out to Robiul for guiding me through the fundamentals of acoustics without ever making me feel like my questions belonged on a “Frequently Asked and Easily Googled” list.

I gratefully acknowledge the Provost’s Awards of Trinity College Dublin for funding my research. Without it, my PhD might still be in the “aspirational” phase — and my diet would certainly have included far fewer noodles.

To my siblings and friends (old and new), thank you for your prayers, encouragement, and unconditional support, even when I vanished for weeks into the academic abyss. May your kindness be returned to you in full, and with interest.

To my mum and dad: thank you for your training, values, and belief that I could reach high places, even when I was still figuring out how to reach the top shelf in the kitchen.

And above all, to God, for the hope of life, the grace to complete this work, and for holding me together when my equations refused to. In thy hands, o Lord!

Contents

Declaration	i
Abstract	ii
Acknowledgements	vii
List of Figures	xxv
List of Tables	xxvii
Nomenclature	xxxix
1 Introduction	1
1.1 Context and motivation	2
1.2 Problem statement	3
1.3 Research aim and objectives	5
1.4 Key definitions and terminology	6
1.5 List of publications	7
1.6 Thesis outline	8
2 Literature Review	10
2.1 Introduction	11
2.2 Acoustic materials	12
2.2.1 Sound absorbing materials	13
2.2.2 Resonant absorbers	14
2.2.3 Classic Helmholtz resonator	14
2.2.4 Perforated panel absorbers	17
2.2.5 Micro-perforated panel absorber	19
2.3 Metamaterials: An overview	20
2.3.1 Photonic crystals	22
2.3.2 Phononic crystals	23

2.3.3	Acoustic metamaterial	24
2.4	Acoustic metasurface	28
2.4.1	Physics of acoustic metasurface	29
2.5	Review of acoustic metasurface designs	34
2.5.1	Acoustic metasurfaces for controlled reflection	34
2.5.2	Acoustic metasurface for controlled transmission	35
2.5.3	Acoustic metasurfaces for controlled absorption	38
2.6	Deep learning optimisation	42
2.6.1	Review of deep learning models in AMM design	43
2.7	Conclusions	48
3	AMM Modelling and Deep Learning Frameworks	51
3.1	Introduction	52
3.2	Absorber description	53
3.2.1	Analytical model	55
3.2.2	Broadband model	58
3.3	Numerical model setup	59
3.3.1	Linearised Navier-Stokes equations	60
3.4	Additive manufacturing	61
3.5	Experimental setup	62
3.6	Deep learning framework	64
3.6.1	Data acquisition	64
3.6.2	Model architecture	67
3.6.3	Data Augmentation	69
3.6.4	Bandwidth modulation	71
3.6.5	Peak scaling	71
3.6.6	Frequency shifting	71
3.6.7	Noise addition	72
3.6.8	Transfer learning	73
3.7	Conclusions	78
4	Model Validation and Parametric Studies	80
4.1	Introduction	81
4.2	Numerical modelling	81

4.2.1	Test case 1	82
4.2.2	Test case 2	83
4.3	Simplified numerical model	85
4.3.1	Effect of the number of absorber units on acoustic response	86
4.4	Experimental validation	90
4.5	General design considerations	92
4.5.1	Sensitivity analysis	93
4.6	Conclusions	95
5	Traditional and Deep Learning Methods in AMM Design	98
5.1	Introduction	99
5.2	Design principle	100
5.3	Traditional optimisation	100
5.3.1	Generalised pattern search	101
5.3.2	Deep learning model	104
5.4	Comparative study	107
5.4.1	Robustness to parametric uncertainty	107
5.4.2	Time complexity	112
5.5	Data quality consideration	117
5.5.1	Peak information data	117
5.5.2	Peak data enriched with bandwidth information	118
5.5.3	Full spectrum data	118
5.6	Conclusions	120
6	Robust Design Considerations	122
6.1	Introduction	123
6.2	Motivation for fabrication-robust design	124
6.3	Data visualisation	124
6.4	Results and discussion	126
6.5	Manufacturing uncertainty consideration	131
6.5.1	v-fold cross-validation strategy	132
6.6	Robustness analysis	134
6.6.1	Robustness to parametric uncertainty	135
6.7	Rapid on-demand design	136

6.8	Conclusions	137
7	Interpretable Inverse Modelling for Geometry Discovery through Design Space Segmentation	139
7.1	Introduction	140
7.2	Data-driven applications in acoustics	141
7.3	Characterisation of the absorber system	142
7.3.1	Design-space mapping of spectral clusters	146
7.4	Discriminative analysis based on ANOVA	148
7.5	Design implications based on cluster characteristics	150
7.5.1	Cluster-specific geometric profiles	150
7.6	Surrogate classifier	151
7.7	Design scenarios	154
7.7.1	Design objective	155
7.8	Geometry retrieval from cluster ID	155
7.9	Practical application to room acoustics	161
7.9.1	Room description and numerical modelling	163
7.9.2	Reverberation time	164
7.9.3	Measurement setup	165
7.9.4	Model validation and results	165
7.10	Conclusions	170
8	Conclusions	172
8.1	General conclusions	173
8.2	Research contribution	174
8.3	Future work	175
	Bibliography	201
A	The AMM Absorber Finetuning Process	202
A.1	Introduction	203
A.2	Initial concept and motivation	203
A.3	Iterative adjustments	203
A.3.1	Cavity geometry: from cylinder to hexagon	203
A.3.2	Top plate perforation: from circle to spiral geometry	204
A.3.3	Manufacturing and analytical simplification	205

A.4	Final design characteristics	205
B	Numerical study	207
B.1	Introduction	208
C	Supplementary Parametric Studies	211
C.1	Introduction	212
C.2	Investigated Parameters	212
D	Linear Acoustic Wave Theory	217
D.1	Theory of acoustic wave propagation	218
D.1.1	Solution of the acoustic wave equation	222
D.1.2	Sound behaviour at medium boundaries	223

List of Figures

1.1	Schematic overview of the robust design workflow, from dataset generation based on the underlying AMM absorber geometry to the final robust design.	3
2.1	Interaction of sound wave with a layer of material. The incident sound energy is decomposed into reflected, dissipated and transmitted sound depending on the type of material used	13
2.2	(a) Schematic of a classic Helmholtz resonator consisting of a neck and a cavity. (b) Equivalent mass-spring model representing the dynamic behavior of the resonator, where the mass corresponds to the air in the neck and the spring to the compressibility of the cavity air	15
2.3	(a) Schematic of a perforated panel absorber consisting of a rigid panel with regularly spaced circular holes, backed by an air cavity. The panel thickness, hole diameter, perforation ratio, and cavity depth determine the resonant frequency and overall acoustic absorption performance of the system, and (b) Illustrates the perforated sheet	18
2.4	An illustration of a single perforation in an MPP where each perforation serves as the neck of a Helmholtz resonator, and a long tube acting as the cavity which is backed by a rigid boundary at $x=d$	20
2.5	A periodic array of copper split-ring resonators on fibre glass circuit board [1].	21
2.6	Sonic crystal: First AMM fabricated [2].	27
2.7	A transmission coefficient curve generated from the first $8 \times 8 \times 8$ sonic crystal [2].	27
2.8	Membrane type AMM [3].	28

2.9	Schematic illustration of generalised Snell's law and acoustic metasurface. The interface is treated with an engineered periodic structure which results in an instantaneous phase shift along the acoustic path, and the generalised Snell's law of reflection and refraction occurs at the interface	30
2.10	Acoustic metasurfaces based on controlling the reflected wave by wavefront manipulation. Figure (a) is the phase shifts of the reflected waves from four different types of labyrinthine units in Figure (b), where Figure (b) is a 3D printed labyrinthine structure with eight labyrinthine units. The label (m, n) indicates that m plates from the top and n plates from the bottom are used for a labyrinthine unit. The eight dots in Figure (a) refer to the chosen eight units shown in B [4]. Figure (d) is a schematic of the Schroeder diffuser. The top and bottom images are the top and 45° angle views o, while the insets is the cross section of unit cells in SD and MSD, with the thicknesses being $\frac{\lambda_0}{20}$ [5]	36
2.11	Metasurface structures of the (a) coiled-up type, (b) Helmholtz type, and (c) membrane type	36
2.12	Fabricated supercell of an acoustic metasurface structure (top image) for anomalous deflection of transmitted acoustic wave [6]. Composed of spatially varied coiling-slit subunits, with a deep subwavelength thickness. By optimizing its microstructures, high performance wavefront steering is exhibited. (b)-(d) Measured amplitude, temporal, and phase field distributions. The green arrows and parallel lines indicate the propagating wavefronts	37
2.13	Fabricated supercell of a tapered labyrinthine metasurface structure (left) and phase changes corresponding to six types of unit cells [7].	38
2.14	Acoustic metasurface composed of a centred perforated plate and a coiled air chamber (left) and the corresponding acoustic response (right) at a resonant frequency of 125 Hz [8].	41
2.15	A multi-coiled metasurface absorber with an embedded aperture (left) and the corresponding acoustic absorption response (right) which achieves a full absorption at a resonant frequency of 50 [Hz] [9].	42
3.1	Schematic overview of the methodology adopted in this chapter, illustrating the complete workflow from AMM concept design through modelling, data-driven learning, and robust design integration.	53

3.2	The unit structure of the absorber: (a) disassembled system comprising a rigid backing, hexagonal boundary and a top perforated plate, (b) parameterised spiral, and (c) assembled system	54
3.3	(a) Schematic of the supercell composed of four inhomogeneous units arranged parallelly, in an edge-to-edge hexagonal tiling. While four units are considered in this thesis, the configuration can be generalised to include more units as required by the application.	59
3.4	Schematic of the two-microphone impedance tube setup used for acoustic absorption measurements. Adapted from Hossain et al. [10].	63
3.5	Photograph of the experimental impedance tube setup used for acoustic absorption measurements in a typical testing session. Key components include the loudspeaker, two calibrated microphones, sample holder, and data acquisition system.	63
3.6	Data distribution scheme used for deep learning training and validation. The diagram illustrates how the dataset is partitioned into training, validation, and test subsets, supporting model development and performance evaluation . . .	65
3.7	Schematic of the proposed SLAE architecture, consisting of an inverse model (encoder) that predicts a 2D latent vector of geometric features from the given spectral features, and a forward model (decoder) that reconstructs the original input features from the latent space. The size-related features (L_s and H) are directly passed into the latent space to support reconstruction. The encoder comprises 4 fully connected layers, including the latent layer, with ReLU activations, while the decoder mirrors this architecture. Both modules incorporate layer normalization (LN).	69

3.8	Illustration of sample data augmentation scenarios. The red curve represents the original (theoretical) spectrum, while the black curves depict augmented spectra generated using various transformation techniques. These transformations include, (a) bandwidth adjustment to reflect potential broadening or narrowing of spectral features, (b) peak scaling to emulate variations in absorption amplitude, (c) horizontal shifts to simulate frequency drift, and (d) noise addition to account for experimental imperfections. Each augmented spectrum demonstrates the ability of the applied methods to introduce realistic variability while preserving key spectral characteristics.	76
3.9	Schematic representation of the transfer learning process. Only the final two layers of the encoder and decoder, highlighted in red, are trainable, while all other layers remain frozen. The network input consists of experimentally measured spectra that reflect inherent manufacturing imperfections.	76
3.10	Overview of the design workflow for the robust AMM design. The pipeline incorporates a pre-trained model within a transfer learning framework, utilising experimentally measured absorption data.	77
4.1	Numerical model setups for (a) a single unit of the absorber with a cylindrical inlet manifold of the same size and cross-section as the experimental tube, and (b) multiple homogeneous units with a hexagonal inlet manifold of the same size and cross-section as each absorber unit	82
4.2	(a) The 3D printing process used to fabricate the meta-absorber, illustrating the printing setup and methodology, (b) the completed fabricated sample showcasing the final design, (c) the experimental test rig, and (d) shows the experimental result compared to the numerical absorption curve.	83
4.3	Experimental and numerical absorption curves of the absorber with parameters $Ls = 50$ [mm], $t = 0.5$ [mm], $p = 7.2$ [mm], $H = 13$ [mm], $tp = 1$ [mm], $r_o = 2.5$ [mm] and $Nt = 4$	84

4.4	Acoustic field distribution of the HRAM at resonance. (a) Instantaneous total acoustic particle velocity magnitude, showing strong flow acceleration at the aperture entrance. (b) Total thermoviscous power dissipation density, indicating that acoustic energy dissipation is concentrated along the channel walls where viscous and thermal boundary-layer effects dominate. The contrast between velocity and dissipation fields highlights that absorption at resonance is governed by thermoviscous losses near solid boundaries.	85
4.5	Equivalent model of the absorber showing: (a) a quarter-section view of the equivalent model comprising a cavity and an equivalent straight channel which includes the perforation. The manifold is highlighted in blue, and (b) a schematic representation of the coiled channel, with the effective length calculation path explicitly marked to demonstrate the acoustic pathway.	87
4.6	Comparison of absorption performance curves obtained from the simplified numerical model (SNM), full-wave (FW) simulations, and experimental measurements. The absorber parameters used in the simulations are $L_s = 35$ [mm], $t = 0.8$ [mm], $p = 2.7$ [mm], $H = 20$ [mm], $r_o = 2.3$ [mm], and $N_t = 3$. The figure highlights the alignment between the simplified model and the full-wave simulations, as well as their agreement with experimental results, demonstrating the validity and accuracy of the equivalent model for predicting absorber performance.	88
4.7	Schematic arrangement of unit cells on 199.5 [mm] by 199.5 [mm] square panels. These are additively manufactured and responses are compared and shown in Fig. 4.8.	89
4.8	Comparison of experimental absorption curves for units 1 through 4 (shown as circular markers) and the numerical absorption curve (depicted by the black solid line). The plots show that absorptivity increases with increase in the number of units printed on the panel and convergence is reached at infinity, corresponding to the numerically stimulated result with infinite surface assumption. Thus, four units are optimal and infinite in this case to realize the true absorption performance of the absorber.	89

4.9	(a) Comparison of the analytical absorption curve with numerical and experimental results. The insets show the CAD designs of the four geometric cases from Table 4.1 that are fabricated.	92
4.10	Influence of size-related parameters and the number of turns on resonance behavior. (a) and (b) are contour plots illustrating the resonance behavior of the absorber at varying values of Ls and H , while keeping other geometric parameters constant: $t_p = 1.5$ [mm], $N_t = 3$, $p = 5$ [mm], and $t = 0.8$ [mm], and (c) illustrates the effect of varying N_t on resonance behavior, with the parameters $t_p = 1.5$ [mm], $p = 4.2$ [mm], $t = 0.8$ [mm], $Ls = 63.5$ [mm], and $H = 15$ [mm], held constant. This highlights the interplay between Ls and H as well as the impact of increasing the channel length through N_t on the resonant frequency.	93
4.11	Absorber responses to variations in (a) p and (b) t . The absorber responses are more sensitive to changes in t than in p , implying that the wall thickness plays a significant role in shaping the acoustic response of the absorber. . . .	94
4.12	The real and imaginary parts of the impedance for different channel widths, p , and spiral wall thickness, t . Zero crossings of the imaginary part indicate resonances, though absorption remains below unity in most cases due to incomplete impedance matching	95
5.1	Architecture of the DNN with layer input and output dimensions and those of the weights and biases. $\mathbf{S}^L$ denotes the number of artificial neurons in layer L , $\forall L = 1, 2, \dots, n$	105
5.2	Training and validation learning curves of the DNN model showing rapid decrease in the loss and close convergence of the training and validation loss, indicating no signs of overfitting	106
5.3	Absorption spectra obtained using the predicted geometric parameters from the DNN (black) and GPS (gray). Both predictions closely align, reproducing the targets ($f_r = 110$ [Hz] and 151 [Hz]) reasonably well.	108

5.4	Histogram summary of robustness metrics computed over 905 MC realisations. Each bar corresponds to one robustness metric, mean peak absorptivity, peak variability, fraction within $\pm 1\sigma$, and fraction outside $\pm 2\sigma$, with error bars showing the standard deviation across runs. The visual representation allows direct comparison between DL and GPS designs across all robustness criteria.	110
5.5	Comparison of robustness between DL and TO designs under FDM manufacturing perturbations. Shaded regions represent $\pm 2\sigma$ prediction bands across the frequency spectrum, and solid black lines indicate mean absorption. The prediction band from GPS is slightly narrower but on the overall, both methods achieve good robustness property.	111
5.6	Comparison of time complexity between TO (GPS) and (b) DNN. The computational cost of TO increases approximately linearly with input size, $O(N)$, whereas the DNN achieves a near-constant inference time, $O(1)$.	114
5.7	Time complexity plots of (a) IDDNN and (b) GPS.	115
5.8	Comparison of DNN performance under different input feature configurations. Three input scenarios are evaluated: (i) peak-based features, (ii) bandwidth-enriched features, and (iii) full-spectrum features. The bandwidth-enriched input achieves the best balance of accuracy and dimensionality, yielding lower MSE and higher R^2 compared to both peak-only and full-spectrum inputs.	120
6.1	Histogram distributions of key spectral parameters in the training dataset. (a) Peak absorptivity showing a strong concentration in the high absorptivity region. (b) Peak frequency exhibiting noticeable non-uniformity and modal clustering. (c) Bandwidth distribution, revealing a strong skew toward narrowband responses.	125
6.2	Learning curves of the SLAE on training and validation datasets. Both training and validation losses converge closely to a minimum, indicating stable learning and no significant overfitting.	126

6.3	Kernel density plots of the errors for outputs p and N_t (top row), and linear regression plots comparing true versus predicted values (bottom row). The vertical dashed lines in the KDE plots indicate the mean and one standard deviation from the mean for each error data. The linear regression plots show the relationship between the true and predicted values for each output. Labels (a), (b), (c), and (d) correspond to each plot as described in the figure.	127
6.4	Predicted absorption spectra from SLAE (black lines) compared to the true spectra (red dots) for 3 randomly generated samples from the analytical model. The plots illustrate the model's ability to accurately reconstruct the spectral features.	128
6.5	Predicted absorption spectra from SLAE (black lines) compared to the true spectra (red dots) for 3 random samples based on finite element simulations. The plots illustrate the model's performance to previously unseen spectral characteristics.	129
6.6	Learning curve of the SLAE trained with a mix of real and augmented data. The plot depicts the training and validation losses over successive epochs, highlighting the model's convergence behavior. The inclusion of augmented data introduces additional variability in the training and validation sets, which can enhance the model's generalization capabilities and mitigate overfitting. The relatively stable validation loss demonstrates the robustness of the model under these conditions.	130
6.7	Comparison of spectra obtained from SLAE-AUG, the unaugmented SLAE, and numerical simulations. The results highlight the improvements in spectral accuracy with augmentation, particularly in regions where the mismatch between the unaugmented model and the numerical response is most pronounced.	131
6.8	Comparison of absorptivity spectra highlighting the impact of transfer learning in capturing manufacturing-induced uncertainties. (a) Analytical and numerical models show clear discrepancies with the experimental spectrum, particularly in peak absorptivity and bandwidth. (b) Incorporation of transfer learning via SLAE-TL substantially reduces this mismatch, yielding predictions that closely follow the experimental curve. This demonstrates the ability of the TL to address the gap between theoretical models and fabricated AMMs.	134

6.9	Comparison of robustness between SLAE and SLAE-TL designs under FDM manufacturing perturbations. Shaded regions represent $\pm 2\sigma$ prediction bands across the frequency spectrum, and solid black lines indicate mean absorption.	135
6.10	Demonstration of the SLAE-AUG network's performance in rapid, on-demand inverse design. The network is provided with user-defined input waveforms representing target spectral behaviors, including (a) rectangular, (b) Gaussian, (c) triangular, and (d) another rectangular profile. For each case, the predicted absorption spectra generated by the network closely follow the desired waveform, illustrating the model's flexibility and accuracy even in the absence of experimentally or numerically derived input spectra.	137
7.1	Cluster analysis of the absorber system using k-means clustering with $k = 5$. The data points represent peak absorptivity and the corresponding frequency extracted from the spectral response. Each cluster groups similar absorber configurations, revealing underlying trends in acoustic performance.	144
7.2	(a) Averaged absorptivity spectra corresponding to each of the five clusters, illustrating distinct spectral behaviors. (b) Histogram showing the number of samples per cluster, with Cluster 4 dominating the distribution, indicating a significant concentration of responses in the low-frequency range.	146
7.3	Boxplot distributions of geometric design parameters across the identified clusters. Each subplot corresponds to a specific parameter: (a) spacing between turns, p , (b) number of turns, N_t , (c) side length of cavity, L_s , and (d) absorber thickness, H . The boxplots reveal the intra-cluster variability and inter-cluster differences in parameter distributions, providing insights into how design parameters are grouped within each cluster. Notably, some parameters exhibit strong clustering trends, such as distinct separation in p and N_t , while others show overlapping ranges, indicating shared design characteristics.	147
7.4	Heatmap of ANOVA F-values showing the statistical significance of each geometric parameter across identified clusters. Higher F-values indicate greater variance between a specific cluster and the remaining dataset for a given parameter, reflecting its discriminative importance in the clustering structure. The color scale highlights the magnitude of these differences.	150

7.5	Optimal search for the number of principal components. The figure illustrates the evaluation process used to determine the ideal number of principal components to retain for dimensionality reduction. By analyzing the variance explained across different numbers of components, the optimal balance between model complexity and accuracy is determined.	152
7.6	Schematic illustration of the Random Forest (RF) classifier architecture. The model consists of an ensemble of $N = 200$ decision trees, each trained on a bootstrapped subset of the training data with randomly selected feature subsets at each split. The RF aggregates the predictions from all constituent trees through majority voting to assign a final class label. This ensemble strategy enhances generalization, reduces overfitting, and improves robustness to noise in high-dimensional input spectra.	153
7.7	Confusion matrix illustrating the classification performance of the surrogate model in predicting geometric cluster labels from absorption spectra. The diagonal entries indicate correctly predicted samples for each cluster, while off-diagonal elements represent misclassifications. High values along the diagonal suggest strong model accuracy, with minimal cross-talk between clusters. This performance supports the utility of the surrogate model in reliably mapping spectral responses to design-relevant geometric groupings.	154
7.8	Modularized pipeline of the segmentation framework. The workflow consists of: (1) applying PCA to compress the spectral data; (2) clustering the reduced representation using k-means to define discrete regions in the design space; (3) training an RF surrogate to predict the cluster label of new spectral inputs; and (4) retrieving m geometry candidates using the nearest-neighbor heuristics within the predicted cluster.	157
7.9	The target absorption spectrum (shown in red circles) is provided as input to the NNH module depicted in Fig. 7.8. The classifier predicts the most appropriate cluster label from the segmented design space, which is then used by the NNH to retrieve candidate geometries. The curves labeled G1 through G5 represent the simulated absorption responses of the retrieved geometries listed in Table 7.2.	159

7.10	Performance of the segmentation framework to varying input types. The predicted acoustic responses are based on analytically simulated geometric parameters retrieved from the NNH module. This allows to validate the accuracy of the retrieved parameters. All cases use scoring weights $\gamma = 1$, $\beta = 0.1$, and $\zeta = 3$, unless otherwise specified.	160
7.11	Illustration of the office space drawing simulated in COMSOL. (a) shows the CAD drawing of the room with a table, book shelve and a drawer, and (b) shows the map of the speaker-microphone positions during measurements. .	164
7.12	Acoustic response of the space measured using a logarithmic sine sweep excitation showing the reverberation time (RT60) and decay characteristics of the space.	166
7.13	Comparison between experimental and numerical reverberation times after correcting for the unknown absorption area in the space.	167
7.14	Comparison between the predicted response and the desired input curve, generated from 4 Gaussians with randomised peak heights, based on the region of interest as observed in the room's reverberation response.	168
7.15	Feature size comparison of the absorber designs.a) a 60 [cm] $\times$ 60 [cm] Spektrum A10 industrail panel used as a benchmark, b) a sample fabricated structure of the AMM absorber	168
7.16	Performance comparison of RT60 curves from both absorbers relative to the baseline.	169
7.17	Relative difference in RT60 from the baseline from both absorbers in the targeted frequency range.	169
A.1	Cavity design shapes	204
A.2	Perforation geometries	205
A.3	Sample 3D-printed AMM units mounted on a 199.5 $\times$ 199.5 [mm] square panel. The design demonstrates the efficient dense packing of hexagonal unit cells with minimal interstitial space between the unit cells. Subfigure (a) shows the external view of the fabricated sample, while subfigure (b) presents a transparent view of (a), allowing visualisation of the internal packing.	206

B.1	Mesh convergence analysis. The top subplot shows the variation in resonance frequency and peak absorption with increasing mesh density (measured by degrees of freedom, DoF). The bottom subplot illustrates the corresponding increase in computational time. The mesh level with DoF = 122,286 provides a good balance between solution accuracy and computational cost.	210
C.1	Variations in acoustic response due to changes in tp . Here, $Ls = 35$ [mm], $Nt = 3$, $p = 2.7$ [mm], $t = 0.8$ [mm], and $H = 15$ [mm], are fixed	212
C.2	Plot showing the peak absorption as a function of r_o and peak frequency, f_r , under variations in $(p + t)$	214
C.3	Plot showing the peak absorption as a function of initial radius, r_o and peak frequency, f_r , under variations in H	216
D.1	Infinitesimal fluid volume element	219
D.2	A harmonic plane wave with incident acoustic pressure, p_i , originates from medium 1 and is normally incident on a planar boundary at $z = 0$. The wave produces a reflected pressure, p_r , in medium 1 and a transmitted pressure, p_t , in medium 2.	223

List of Tables

3.1	Geometric design parameters defining the unit cell of the metasurface acoustic absorber. These parameters control/influence the absorption behaviour of the absorber.	55
3.2	Fluid simulation parameters.	61
3.3	Parameter ranges for deep learning, where $t = 0.8$ [mm] and $tp = 1.5$ [mm] are fixed	66
3.4	Summary of the proposed SLAE architecture, showing the encoder-decoder layer types, output shapes, and number of trainable parameters	70
4.1	Sample geometric parameters of the simulated absorber with $tp = 1$ [mm] in all cases	91
4.2	Air simulation parameters used in all models	91
5.1	Quantitative assessment of DNN performance on the predicted outputs. . . .	106
5.2	Comparison of predicted output values for geometric parameters from the DNN and GPS models.	107
5.3	Robustness comparison between DL and GPS over 5 random seeds. Metrics are reported as mean $\pm$ standard deviation.	110
5.4	Performance of the DNN under different input feature configurations.	119
6.1	Summary of the evaluation metrics used to assess the network's performance.	128
6.2	Sample Geometric parameters of the simulated absorber with $t = 0.8$ [mm] and $tp = 1.5$ [mm] in all cases	128
6.3	Summary of the evaluation metrics for the augmented dataset scenario. The results demonstrate good predictive performance of the model under this condition, with R^2 values exceeding 90% and sufficiently low MSE values	130

6.4	Key uncertainty sources and their ranges for PLA and Prusa i3 MK3S+ machine	132
7.1	Cluster-wise summary of dominant geometric features and their acoustic im- plications	151
7.2	Geometric parameter values of the 5 retrieved geometries corresponding to the target spectrum within the predicted Cluster 2.	158
7.3	Comparison between deep learning and segmentation-based methods for in- verse modeling.	161
7.4	Absorption coefficients of the surface materials in the investigated office space.	164
B.1	Mesh convergence analysis	209
C.1	Changes in r_o due to changes in $(p + t)$, keeping other parameters constant, i.e., $Ls = 35$ [mm], $Nt = 3$, $t = 0.8$ [mm], $H = 15$ [mm], and $tp = 1.5$ [mm] .	213
C.2	Changes in r_o due to changes in $(p + t)$, keeping other parameters constant, i.e., $Ls = 35$ [mm], $Nt = 3$, $t = 0.8$ [mm], $p = 3.5$ [mm], and $tp = 1.5$ [mm] .	215

Nomenclature

Acronyms

AI Artificial Intelligence

AM Additive Manufacturing

AMM Acoustic Metamaterial

ANC Active Noise Control

BEM Boundary Element Method

CAD Computer Aided Design

CGAN Conditional Generative Adversarial Network

CNN Convolutional Neural Network

DL Deep Learning

DNN Deep Neural Network

EMM Electromagnetic Metamaterial

FDM Fused Deposition Modelling

FEM Finite Element Modelling

FWHM Full Width at Half Maximum

GAN Generative Adversarial Network

GPS Generalised Pattern Search

HR Helmholtz Resonator

HRAM Hybrid Resonant Acoustic Meta-Absorber

HVAC Heating, Ventilation, and Air Conditioning

JCAPL Johnson-Champeaux-Allard-Pride-Lafarge

MC Monte Carlo

MPP Micro-perforated Plate

MSE Mean Squared Error

NRC Noise Reduction Coefficient

PCA Principal Component Analysis

PLA Polylactic Acid

PML Perfectly Matched Layer

PNC Passive Noise Control

RNN Recurrent Neural Network

SLAE Structured Latent Autoencoder

SNM Simplified Numerical Model

SPL Sound Pressure Level

t – SNE t-Distributed Stochastic Neighbor Embedding

TL Transfer Learning

TMF Transfer Matrix Formulation

TO Traditional Optimisation

UMAP Uniform Manifold Approximation and Projection

VAE Variational Autoencoder

Greek Symbols

α_p Peak absorptivity

-

α	Absorption coefficient	—
δ	Viscous boundary layer thickness	m
λ	Wavelength of sound	m
ω	Frequency	$rad s^{-1}$
$\phi(x)$	Phase shift function	rad
ϕ	Porosity/Open area ratio	—
ρ	Mass density	$kg m^{-3}$
ρ_0	Equilibrium density	$kg m^{-3}$
θ_i	Incident angle	rad
θ_r	Reflected angle	rad
θ_t	Transmitted angle	rad

Roman Letters

B	Adiabatic bulk modulus	Pa
c	Speed of Sound	$m s^{-1}$
d	Impedance tube diameter	mm
f_L	Lower band-edge frequency	Hz
f_r	Resonant frequency	Hz
f_u	Upper band-edge frequency	Hz
H_{12}	Complex transfer function	—
k	Wave number	m^{-1}
p	Acoustic pressure	Pa
p_+	Acoustic pressure of forward propagating wave	Pa
p_-	Acoustic pressure of backward propagating wave	Pa

p_0	Equilibrium pressure	Pa
R	Reflection coefficient	—
s	Spacing between microphones	mm
T_j	Transfer matrix of region, j	—
v_0	Mean velocity	ms^{-1}
Z	Normal acoustic impedance	$kgm^{-2}s^{-1}$

Chapter 1

Introduction

1.1 Context and motivation

Noise, unlike desirable sounds such as music or speech, is generally intrusive or harmful. Its sources are diverse, including transportation systems, industrial processes, and urban activities. Chronic exposure to environmental noise has been linked to cardiovascular diseases, sleep disturbance, cognitive impairment, and psychological stress [11, 12]. Indoor noise, in particular, can disrupt learning, communication, and productivity in workplaces, schools, and healthcare facilities [13]. Controlling noise in these environments is therefore critical for both health and functional performance.

The effectiveness of conventional acoustic materials is often limited, especially at low frequencies. Porous and fibrous materials rely on viscous and thermal losses for absorption, which are most effective when the material thickness approaches a quarter of the target wavelength. This makes low-frequency attenuation impractical due to the large thickness required [14, 15]. Resonant absorbers, such as Helmholtz resonators and micro-perforated panels, can target specific low frequencies but typically operate over narrow bands and may become bulky when designed for extreme low-frequency attenuation.

Acoustic metamaterials (AMMs) provide a promising alternative, enabling subwavelength control of sound, tunable absorption, and lightweight structures [2, 16]. By manipulating the geometry of the material’s building blocks—such as coiled channels, labyrinthine paths, and embedded resonators—AMMs can achieve enhanced low-frequency attenuation in compact volumes where traditional absorbers would require prohibitive thickness. These materials enable applications such as low-frequency room acoustic control, compact HVAC noise reduction, and targeted noise mitigation in industrial and transportation contexts.

Despite these advantages, designing AMMs is challenging. The acoustic performance can be highly sensitive to geometric parameters, manufacturing imperfections, and modeling assumptions. Traditional design approaches, relying on trial-and-error experimentation or computationally expensive simulations, are time-consuming and often lack interpretability. Additive manufacturing (AM) enables the production of complex geometries, but introduces variability that can degrade performance if not properly accounted for [17, 18].

This thesis addresses these challenges by developing a data-driven framework for the design of AMM absorbers that is robust, interpretable, and predictive. The framework integrates analytical modeling, high-fidelity simulations, and deep learning techniques to enable both forward and inverse design, while incorporating manufacturing tolerances and real-world

uncertainties. By linking geometric features to acoustic performance, the approach not only improves efficiency and robustness but also provides scientific insight into design principles. Figure 1.1 illustrates the general workflow for the robust AMM design adopted in this thesis.




Figure 1.1: Schematic overview of the robust design workflow, from dataset generation based on the underlying AMM absorber geometry to the final robust design.

1.2 Problem statement

AMMs have garnered significant interest in recent years due to their ability to manipulate sound waves in unconventional ways, which opens up new possibilities for applications in areas such as noise control, acoustic focusing, vibration isolation, and soundproofing. These materials have the potential to achieve effects that traditional materials cannot, such as negative refraction, acoustic cloaking, and sound wave steering. However, the design and optimization of AMMs remain complex and resource-intensive processes, presenting significant challenges for researchers and engineers working in the field. The traditional design approach for AMMs relies heavily on trial-and-error experimentation, a method that can be both time-consuming and inefficient. Designers often have to explore a vast design space of possible geometries and materials, testing various configurations to find those that meet the desired acoustic properties. This process can involve a significant amount of manual tuning and iteration, which is

not only slow but also prone to sub-optimal outcomes due to the reliance on intuition and experience rather than systematic exploration of the design space. One of the key challenges in AMM design is that the desired acoustic behaviour is highly dependent on the geometry of the material. These properties, such as sound absorption, transmission, and reflection, can be deeply influenced by factors like the shape, size, orientation, and spacing of the constituent elements. Since the relationship between the material’s geometry and its acoustic behavior is often nonlinear, predicting the performance of a given configuration can be difficult without extensive simulation. While traditional optimisation methods, such as gradient-based techniques or global optimizers like genetic algorithms, can be effective, they are often laborious and computationally expensive. These methods typically evaluate a vast number of potential solutions to find an optimum, which can be prohibitively costly in terms of time and resources. Additionally, gradient-based methods may struggle with local minima and non-smooth design spaces, often leading to sub-optimal results. In light of these challenges, the need for alternative, more efficient design methodologies is evident. The traditional methods are limited in their ability to handle the high dimensionality and complexity of the design problem. This calls for the exploration of methods that are more efficient, providing solutions that are both optimised and computationally feasible. Secondly, additive manufacturing (AM), also known as 3D printing, has emerged as a revolutionary method for fabricating complex material structures, offering the ability to design and produce AMMs with unprecedented geometric freedom. While state-of-the-art artificial intelligence (AI) (machine/deep learning) methods have revolutionised the design landscape, the process hitherto does not account for the real-world imperfections that arise in the manufacturing process. AM is prone to variations in material properties, geometric tolerances, and structural inconsistencies, which can significantly impact the performance of the designed AMMs [17–21]. The lack of robustness in the modern use of AI (and traditional) design methods limits the practical application of AMMs, as they may not perform as expected when subjected to manufacturing variations or environmental factors. Thus, incorporating robust design principles that account for these imperfections is crucial for achieving reliable, high-performance AMMs. Thirdly, state-of-the-art deep learning methods for AMM inverse design often operate as black boxes, offering little scientific intuition about why a particular geometry is predicted. This lack of interpretability limits their usefulness for designers, who require not just accurate predictions but also insight into the geometric–acoustic relationships that drive performance. An

interpretable framework allows researchers to trace spectral features back to dominant geometric parameters, providing physical understanding, design control, and clearer pathways for corrective tuning when experimental responses deviate from simulations. By linking performance outcomes with geometric trends, interpretability strengthens both the reliability and the practical adoption of data-driven AMM design strategies.

To this end, there is a pressing need for efficient and robust design pipelines that explicitly incorporate fabrication tolerances while also offering interpretable connections between geometry and acoustic performance. Such methodologies would ensure that AMMs perform optimally not only under idealized simulations but also in real-world applications where imperfections are unavoidable. This thesis therefore addresses the dual challenge of robustness and interpretability in AMM design. By developing data-driven frameworks that are both reliable and transparent, this work advances the practicality of machine-learning-based metamaterial design, bridging the gap between theoretical predictions and experimental realizations.

1.3 Research aim and objectives

This thesis addresses the dual challenges of robustness and interpretability in AMM design. The overarching aim is to develop data-driven frameworks that enable the efficient design of high-performance, fabrication-robust, and interpretable AMMs for noise control. The specific objectives are:

- To conceptualise and design a suitable AMM absorber structure, referred to as hybrid resonant acoustic meta-absorber (HRAM) that is optimisable, easy to manufacture, simple in structure, yet sufficiently complex for the the DL studies.
- To validate the proposed HRAM design numerically, analytically, and experimentally.
- To collect a representative dataset of the validated HRAM for data-driven studies, and to evaluate the use of DL models over traditional optimisation methods based on robustness and time complexity analyses.
- To select and train a suitable DL architecture for both forward and inverse design, enabling prediction of acoustic responses from HRAM geometry and, conversely, HRAM geometric parameters from target responses, while incorporating manufacturing imper-

fections and compensating for model imperfections through data augmentation techniques.

- To integrate pre-trained DL model into a transfer learning (TL) framework, leveraging experimental data that incorporate AM defects.
- To develop interpretability tools that provide physical insights into the design spaces, revealing key design features and decision-making criteria that influence performance.
- To deploy the robust DL HRAM model to design frequency-targeted acoustic panels for practical room acoustic application.

1.4 Key definitions and terminology

For clarity, the following terms are used throughout this thesis with specific meanings:

- **Acoustic Metamaterial (AMM):** An engineered structure whose acoustic behavior is governed primarily by its geometry rather than its base material properties, enabling subwavelength control of sound.
- **Robust Design:** A design approach that explicitly accounts for uncertainties—such as manufacturing tolerances, modeling assumptions, and experimental variability—to ensure consistent performance under real-world conditions.
- **Interpretability:** The ability to extract physically meaningful relationships between design parameters and acoustic performance from a model, allowing designers to understand why a particular geometry works.
- **Deep Learning (DL):** A class of machine learning methods based on multi-layer neural networks capable of learning complex, nonlinear mappings between high-dimensional inputs and outputs.
- **Inverse Design:** A design paradigm in which desired acoustic performance is prescribed as input, and the corresponding geometric parameters are predicted as output.
- **Transfer Learning (TL):** A machine learning strategy where a model trained on a large, simulated dataset is fine-tuned using limited experimental data to improve real-world predictive accuracy.

These definitions provide a consistent conceptual framework for the methods and analyses developed in this thesis.

1.5 List of publications

This thesis is informed by several journal and conference publications authored or co-authored by the candidate. For clarity, these publications are categorised below based on their relevance to the core contributions of the thesis and the candidate's role in each work.

Journal papers forming the core of this thesis:

The following journal papers directly underpin the methodology, analysis, and contributions presented in this thesis. The candidate was the lead author and primary contributor to the conceptual development, modeling, data analysis, and manuscript preparation.

- [1] Ogun, O., Rice, H., & Kennedy, J. (2025). Bridging model and experiment in the design and validation of a sub-wavelength acoustic metamaterial. *Journal of Applied Physics*, 137(19).
- [2] Ogun, O., & Kennedy, J. (2025). Design-space segmentation of an acoustic metamaterial for interpretability and insight. *Journal of Sound and Vibration*, 119443.
- [3] O. Ogun and J. Kennedy, Modeling Acoustic Metamaterials with Manufacturing Uncertainties Using Deep Learning. *Journal of Sound and Vibration*

Conference publications arising from this thesis:

The following conference papers disseminate intermediate results, methodological comparisons, and application-focused extensions of the research presented in this thesis.

- [1] Ogun, O., & Kennedy, J. (2022, September). Noise Attenuation Through Optimised Acoustic Metamaterials: A Low Form Factor Design for Targeted Noise Reduction. In 2022 28th International Workshop on Thermal Investigations of ICs and Systems (THERMINIC) (pp. 1-4). IEEE.
- [3] Ogun, O., & Kennedy, J. (2024, October). Comparison of traditional and deep learning optimisation for the design of acoustic metamaterials. In INTER-NOISE and NOISE-CON Congress and Conference Proceedings (Vol. 270, No. 10, pp. 1352-1362). Institute of Noise Control Engineering.

[4] Ogun, O., Rice, H. J., & Kennedy, J. The influence of dataset generation on the performance of deep learning for AMM design. ISMA2024 - International Conference on Noise and Vibration Engineering, Leuven.

[5] Ogun, O., & Kennedy, J. (2025). Deep learning-aided design of a broadband low-frequency meta-absorber for room acoustic applications. In Proceedings of Euronoise 2025 (Spain).

Related collaborative journal publications:

The following journal/conference papers were collaborative works led by other researchers. While thematically related to metamaterials and inverse design, they are not directly examined in this thesis and are included for completeness and contextual relevance only.

[1] Muhammad, Ogun O, & Kennedy J. Inverse design of a topological phononic beam with interface modes. *Journal of Physics D: Applied Physics*. 2022 Nov 15;55(1):015106.

[2] Muhammad, Kennedy J, & Ogun O. Design and fabrication of 3D-printed composite metastructure with subwavelength and ultrawide bandgaps. *New Journal of Physics*. 2023 May 18;25(5):053015.

[3] Wasala, S., Zhao, W., Ogun, O., Stevens, L., Sosseh, R., Kennedy, J., & Persoons, T. (2023, November). Acoustic metamaterials for electronics cooling fan noise reduction. In INTER-NOISE and NOISE-CON congress and conference proceedings (Vol. 268, No. 5, pp. 3145-3153). Institute of Noise Control Engineering.

1.6 Thesis outline

This thesis is organised into eight chapters, each building upon the previous to progressively address the overarching aim of developing interpretable and robust frameworks for acoustic metamaterial (AMM) design using data-driven and deep learning methods.

Chapter 2 presents a comprehensive review of the literature. It discusses the impact of noise on human health and quality of life, and an in-depth review of AMMs, highlighting their unique capabilities, design challenges, and limitations. Finally, a review on the use of state-of-the-art deep learning methods as a modeling tool for efficient design of AMM, including the identification of existing gaps that motivate the objectives of this thesis.

Chapter 3 presents a detailed methodological procedure of the approaches used to achieve the thesis objectives. The development of an AMM structure that is easy to fabricate and geometrically rich for optimisation purposes, yet sufficiently complex to capture the various mismatch scenarios that are commonly encountered in AMM modeling, and perhaps difficult to model analytically, is developed. Furthermore, numerical modelling procedures, experimental setups, and additive manufacturing workflows are presented in detail. The chapter also introduces the proposed DL framework, including dataset generation, data augmentation strategies, and a transfer learning pipeline for robust model adaptation.

Chapter 4 provides a comprehensive validation of the proposed HRAM, using finite element modeling (FEM), analytical method and experimental testing using the two-microphone impedance tube method. Parametric studies to investigate the impact and the sensitivities of the geometric parameters on the acoustic response, are also presented.

Chapter 5 provides a comparative study between traditional optimisation and deep learning methods using different metrics, such as robustness and time complexity. It aims to investigate the benefits of deep learning in AMM research and to highlight the areas of strength and situations where either method may be preferred over the other. The chapter also investigates the impact of input data configurations on the performance of DL models.

Chapter 6 develops and tests the robust design framework. Building on the autoencoder architecture described earlier, the chapter implements augmentation and transfer learning strategies to incorporate manufacturing imperfections and model inaccuracies. The resulting robust model is also implemented in an on-demand case scenario, where simulated responses are not available as input to the DL model.

Chapter 7 introduces a segmentation-based framework for interpretable AMM design. By partitioning the design space into meaningful clusters, this methodology provides transparency and scientific insight into absorber behaviour, complementing the predictive power of DL with interpretability. The segmentation-based framework is also applied here to design HRAM panels for room acoustic application.

Finally, Chapter 8 synthesises the findings of the thesis, reflecting on how the objectives were achieved and outlining the contributions to the field. The chapter also proposes avenues for future research, including opportunities for generalising the frameworks developed here to broader classes of metamaterials and practical multi-functional applications.

Chapter 2

Literature Review

2.1 Introduction

The detrimental impact of noise includes hearing loss, stress, sleep disturbances, and cognitive impairments, particularly in sensitive populations such as children and the elderly [22–24]. As such, noise control has become an essential aspect of engineering, with both regulatory frameworks and technological advancements seeking to reduce noise levels in various environments [25, 26].

The control of sound, whether to suppress unwanted noise, enhance speech intelligibility, or shape acoustic environments, has long been a central concern in fields such as architectural acoustics, transportation, industrial design, and consumer electronics. Over the decades, a variety of traditional materials and techniques have been developed to address these challenges, relying on well-established physical principles of sound absorption, reflection, transmission, and scattering. Materials such as porous absorbers, panel resonators, Helmholtz resonators, and micro-perforated panels have formed the foundation of conventional noise control strategies. These systems offer practical solutions across many frequency ranges and applications, particularly in mid-to-high frequency domains. However, traditional acoustic materials and designs often exhibit key limitations, especially in achieving high-efficiency performance at low frequencies, where the wavelength of sound is long and requires correspondingly bulky treatments. Additionally, these materials typically lack the ability to dynamically adapt to different acoustic environments or to provide precise, frequency-selective control over sound propagation. In recent years, the emergence of acoustic metamaterials (AMMs) has opened new pathways for overcoming the constraints inherent in traditional noise control systems. Inspired by the field of electromagnetic metamaterials (EMMs), AMMs are artificially structured materials engineered to exhibit properties not found in nature—such as negative effective mass density or modulus, allowing for unprecedented manipulation of sound waves at sub-wavelength scales. These capabilities have led to innovative devices such as compact sound absorbers, waveguides, lenses, and acoustic cloaks. Building on these advances, acoustic metasurfaces have further pushed the boundaries of wave control by enabling tailored reflection, phase shifting, and redirection of sound using ultrathin, surface-based structures. This shift from bulk materials to functional surfaces has opened up possibilities for low-profile, lightweight, and frequency-tunable acoustic devices. Noise control can be broadly divided into two categories: passive and active [27–32]. Passive noise control (PNC) primarily involves the use of materials and structural modifications to absorb, reflect, or deflect

sound waves. Active noise control (ANC), on the other hand, uses electronic systems to produce sound waves that cancel out unwanted noise. This thesis, however, focuses on PNC, and this chapter provides a comprehensive review of the state-of-the-art in traditional acoustic materials and metamaterials.

This chapter is organised as follows. Section 2.2 reviews conventional noise control strategies, beginning with traditional approaches such as barriers, enclosures, and porous absorbers, before introducing resonant absorbers as a bridge to modern AMM concepts. Section 2.3 introduces the concept of metamaterials, outlining their evolution and foundational principles in the context of acoustics. Section 2.4 discusses acoustic metasurfaces and their underlying physics, with emphasis on the generalised Snell’s law. A review of acoustic metasurface designs targeting reflection, transmission, and absorption objectives is presented in Section 2.5. Section 2.6 reviews the use of deep learning methods in AMM design. Finally, conclusions are drawn in Section 2.7.

2.2 Acoustic materials

Acoustic materials are designed to reduce noise by controlling, directing, and manipulating sound waves. Traditional materials have long been used for sound attenuation in various applications. These materials are typically categorised based on their acoustic properties such as absorption, reflection, and transmission. Thus, they can be classified into enclosures, absorbers and barriers. If a layer of material is placed in the path of a propagating sound wave as shown in Fig. 2.1, and it is desired to prevent sound from being transmitted to medium 2, then the material should reflect or dissipate more but transmit less. Barriers and enclosures come under this category. Similarly, to control the noise in medium 1, along with medium 2, the material should reflect less, transmit less, and dissipate more. Absorbers come under this category. These traditional approaches become impractical for low-frequency sound, as they require excessively large material thicknesses to achieve sufficient attenuation, and lack tunability to target specific frequencies. To address this limitation, resonant absorbers have been introduced. Unlike porous absorbers, resonant devices provide strong, narrowband absorption at specific, tunable frequencies, making them particularly suitable for low-frequency noise control.




Figure 2.1: Interaction of sound wave with a layer of material. The incident sound energy is decomposed into reflected, dissipated and transmitted sound depending on the type of material used

2.2.1 Sound absorbing materials

Absorptive materials are designed to reduce the reflection and attenuate airborne sound energy by converting it into other forms of energy, typically heat [33], through viscous and thermal dissipation mechanisms. Commonly used absorptive materials include foams and fibrous materials, such as polyurethane foam, mineral wool, fiberglass, and cotton fibers. They are widely used in architectural acoustics, vehicle cabins, industrial enclosures, sound-proofing applications, etc., due to their high absorption efficiency across a broad frequency range, and increasingly in the design of AMMs. Sound absorbers can be categorised into porous and resonant absorbers. The former absorb sound via viscous losses in interconnected pores, while the latter use a neck and cavity to create a localised resonance that traps and dissipates sound energy [14, 34]. The performance of sound absorbers is evaluated using the sound absorption coefficient, $\alpha(f)$, or noise reduction coefficient, NRC. α is the fraction of the incident sound energy absorbed by the material, and is a function of frequency. Thus,

$$\alpha(f) = \frac{E_{\text{absorbed}}}{E_{\text{incident}}} = 1 - |R(f)|^2, \quad (2.1)$$

where $0 \leq \alpha(f) \leq 1$. NRC is the average of the measured $\alpha(f)$ for 250 [Hz], 500 [Hz], 1000 [Hz] and 2000 [Hz] octave bands rounded off to the nearest multiple of 0.05. Thus,

$$\text{NRC} = \frac{\alpha_{250[\text{Hz}]} + \alpha_{500[\text{Hz}]} + \alpha_{1000[\text{Hz}]} + \alpha_{2000[\text{Hz}]}}{4}. \quad (2.2)$$

This is an approximate and quick single value metric for material selection, but inappropriate in very low and high frequencies. Thus, sound absorption coefficient remains the standard metric to measure the absorptivity of a material

2.2.2 Resonant absorbers

Resonant sound absorbers of the Helmholtz resonator-type have long been employed across a variety of applications due to their effective performance at low frequencies. These systems are typically built with a perforated outer panel, beneath which lies a porous material backed by a rigid wall. In this configuration, the air contained within the perforations or slits acts as the primary oscillating mass, as the panel itself is generally quite rigid. The restoring force—or stiffness—originates from the enclosed air volume behind the panel, while the porous material contributes to damping through viscous dissipation. Despite their traditional design, these resonant absorbers continue to receive considerable interest in current research. Much of this renewed focus is driven by advancements in our understanding of porous materials and their interaction with acoustic waves [35]. Resonant absorbers can be classified into special categories based on the Helmholtz principle: classic Helmholtz resonator (HR), perforated panel absorber (PPA) and micro-perforated panel absorber (MPP).

2.2.3 Classic Helmholtz resonator

The Helmholtz resonator, as shown in Fig. 2.2(a), consists of an acoustical cavity contained by rigid walls and connected to the exterior by a small opening called the neck. It is a special type of air-spring oscillator. Here, the dimension of the resonator is required to be much smaller than the wavelength to be absorbed, i.e., $V \ll \lambda^3$, where V is the volume of the cavity. The equivalent mass-spring analogy is shown in Fig. 2.2(b). When an incident wave impinges on the neck, the air molecules in it start to vibrate back and forth, while the air inside the cavity provides the restoring force. The air in the neck constitute the mass of the system, while the air in the cavity constitute the spring element. When an incident acoustic wave impinges on the neck of the resonator, the air in the neck oscillates back and forth due to the pressure variation, acting as an inertial mass. Meanwhile, the air trapped in the cavity compresses and decompresses in response, serving as the spring in the system. This behavior is analogous to a mass-spring system, as shown in Fig. 2.2(b), where the resonance

occurs when the mass and spring are in dynamic equilibrium. At resonance, the system exhibits maximum energy absorption. To obtain the fundamental (or natural) frequency of the resonator, it is assumed that the volume, V , and pressure, p , inside the cavity remains spatially uniform and varies only with time.




Figure 2.2: (a) Schematic of a classic Helmholtz resonator consisting of a neck and a cavity. (b) Equivalent mass-spring model representing the dynamic behavior of the resonator, where the mass corresponds to the air in the neck and the spring to the compressibility of the cavity air

Assume that the acoustic pressure inside the neck and cavity is denoted $p_n(x, t)$ and $p_c(t)$, respectively. As air oscillate in and out along the neck, the acoustic pressure causes air inside the cavity to compress and expand. If the compression/expansion is assumed to be adiabatic, and the fluctuations are again sufficiently small, so that $\rho - \rho_0 = \rho_0 = \rho_c$, then according to Eq. (D.7), the pressure inside the cavity is derived as

$$p_c(t) = B \frac{(\rho - \rho_0)}{\rho_0}, \quad c^2 = \frac{B}{\rho}. \quad (2.3)$$

$$p_c(t) = c^2(\rho - \rho_0) = c^2 \rho_c. \quad (2.4)$$

As mass flows into the cavity, the density increases, and vice versa.

$$\frac{d\rho_c}{dt} = \frac{\dot{m}}{V} = \rho_0 = \frac{1}{V} \int_0^t \dot{m} dt, \quad (2.5)$$

where $\dot{m}$ is the mass flow into the cavity. Substituting Eq. (2.4) into Eq. (2.5), we have that

$$p_c(t) = \frac{c^2}{V} \int_0^t \dot{m} dt. \quad (2.6)$$

Applying Newton's second law to the air movement in the neck, and noting that the pressure gradient generated throughout the neck length is opposite to the direction of the velocity since air will always move from a zone of high pressure to a zone of low pressure, hence, the negative pressure gradient constitutes the force that generates the air flow.

$$\rho_0 \frac{dv}{dt} = -\frac{\partial p_n}{\partial x} - F, \quad (2.7)$$

where F is the resistive force due to viscosity and any other resistive forces. From the theory of viscosity and resistive forces, $F \propto v$, and $F \propto \frac{1}{L}$, so that $F = \frac{R_c v}{L}$, where R_c is the constant of proportionality. Now, assuming harmonic solutions in the neck, p_n and $v(t)$ can be written as

$$p_n(t) = p_n e^{j\omega t}; \quad v(t) = v e^{j\omega t}. \quad (2.8)$$

Upon differentiating Eq. (2.8) w.r.t t ,

$$\frac{\partial p_n}{\partial t} = j\omega p_n; \quad \frac{\partial v_n}{\partial t} = j\omega v. \quad (2.9)$$

Equation (2.7) can subsequently be re-written as

$$j\omega \rho_0 v + \frac{R_c v}{L} = -\frac{\partial p_n}{\partial x}. \quad (2.10)$$

Upon integrating Eq. (2.10) w.r.t. x , from 0 to L , we have that

$$[p_n]_0^L = p_n - p_c = j\omega \rho_0 v L + R_c v. \quad (2.11)$$

Differentiating Eq. (2.11) w.r.t. t , we have

$$j\omega p_n = \frac{dp_c}{dt} - \omega^2 \rho_0 v L + j\omega R_c v. \quad (2.12)$$

Differentiating Eq. (2.5) w.r.t t , and substituting the result into Eq. (2.12), the pressure, p_n , can be written as

$$p_n = -j \frac{c^2 \dot{m}}{\omega V} + j \omega \rho_0 v L + R_c v. \quad (2.13)$$

From Eq. (2.5), $\dot{m} = \rho_0 S v$, where S is the surface area of the neck. Substituting this into Eq. (2.13), p_n becomes

$$p_n = R_c v - j \rho_0 c \left[\frac{c S}{\omega v} - \frac{\omega L}{c} \right]. \quad (2.14)$$

The acoustic impedance of the resonator can thus be written as

$$Z = \frac{p_n}{v} = R_c - j \rho_0 c \left[\frac{c S}{\omega v} - \frac{\omega L}{c} \right]. \quad (2.15)$$

The resonant frequency, f_r , of the resonator is obtained when the resistance offered by the system to the flow of sound is minimum. Thus, Z is minimum at resonance so that

$$\rho_0 c \left[\frac{c S}{\omega v} - \frac{\omega L}{c} \right] = 0. \quad (2.16)$$

Solving for the resonant frequency, f_r , yields

$$\omega_r = c \sqrt{\frac{S}{V L}} \Rightarrow f_r = \frac{c}{2\pi} \sqrt{\frac{S}{V L}} \quad (2.17)$$

Due to geometric discontinuity at $x=0$ and $x=L$, a correction factor is normally included in the neck length at both ends. See Crocker [36] for different end correction factors depending on the neck geometry. For a circular aperture, a factor of 0.85 is used to correct for both ends of the neck.

$$f_r = \frac{c}{2\pi} \sqrt{\frac{S}{V(L + 0.85d)}} \quad (2.18)$$

2.2.4 Perforated panel absorbers

As an enhancement to the Helmholtz resonator, perforated panel absorbers are a class of acoustic treatments composed of a thin, rigid sheet—typically made of metal, wood, or high-density polymer, with an array of evenly spaced circular or slotted perforations. This perforated panel is mounted over an enclosed air cavity, which is backed by a rigid, reflective surface, as illustrated in Fig. 2.3(a). The acoustic behavior of this system can be modeled analogously to a mass-spring-damper system. In this analogy, the air in the perforations acts as discrete oscillating masses, while the enclosed air volume in the backing cavity functions

as a compressible spring due to its bulk modulus. The damping, which accounts for energy dissipation, can be introduced either by viscous losses at the perforation edges or by adding porous materials within the cavity.

When an incident acoustic wave impinges on the surface of the panel, it induces pressure fluctuations that drive the oscillation of air plugs within the perforations. This motion compresses and decompresses the air in the cavity, thereby storing and dissipating acoustic energy. The system exhibits a resonant frequency that is determined primarily by the geometry of the perforations (e.g., hole diameter, spacing, and panel thickness) and the depth of the backing cavity [37–40]. At resonance, the system achieves maximum acoustic absorption, making it particularly effective in narrow frequency bands. Perforated panel absorbers are widely used in architectural acoustics, HVAC systems, and noise control applications, where passive, frequency-selective absorption is required.




Figure 2.3: (a) Schematic of a perforated panel absorber consisting of a rigid panel with regularly spaced circular holes, backed by an air cavity. The panel thickness, hole diameter, perforation ratio, and cavity depth determine the resonant frequency and overall acoustic absorption performance of the system, and (b) Illustrates the perforated sheet

The porosity or open area ratio, σ , of a perforated panel can be thought of as effective hole area per unit area of the panel. For evenly spaced circular perforations,

$$\sigma = \frac{\pi d_h^2}{4r_p^2}, \quad (2.19)$$

where r_p is the inter-centre spacing between the holes, and d_h is the diameter of the holes. Equation (2.19) is based on the assumption that the rectangular region marked in red in Fig. 2.3(b) represents a single repeating unit, so that the repetition of this unit makes up the panel, and σ is calculated for a single unit. Now, every perforation can be assumed to be a Helmholtz resonator, so that each would have the same fundamental resonant frequency similar to that of Helmholtz resonator in Eq. (2.18). Thus,

$$f_p = \frac{c}{2\pi} \sqrt{\frac{\pi d_h^2}{4V(L + 0.85d_h)}}, \quad (2.20)$$

where L in this case is the thickness of the perforated sheet or neck length, and $V = r_p^2 h_p$, is the volume of the backing cavity. Again, this expression is derived using the classical Helmholtz resonator model, where the perforations act as acoustic necks and the backing cavity provides the compliance (see for example, [41], for more information). Absorption is maximum when the air gap between the panel and rigid wall is $h_p = \lambda/4$, and minimum when $h_p = \lambda/2$.

2.2.5 Micro-perforated panel absorber

Both Helmholtz resonator and perforated panel absorber have limitations in the absorption peak or maximum absorption they can achieve, and more so, a broadband absorption is not possible. Micro-perforated panel absorbers are thus designed to address these limitations. In contrast to conventional perforated panels, which typically feature perforations on the millimeter to centimeter scale and offer minimal intrinsic acoustic resistance, micro-perforated panels (MPPs) utilise much smaller holes—typically ranging from 0.5 to 1 [mm] in diameter. These fine perforations are designed to independently generate sufficient acoustic resistance and exhibit low mass reactance, enabling effective broadband sound absorption without the need for supplementary fibrous or porous materials size [38]. MPP consists of a thin sheet of material with numerous micro holes cut out on it and the sheet backed by an air cavity followed by a rigid backing. It works exactly like a Helmholtz resonator/perforated panel already discussed. When the holes in an MPP become too small (less than 1 [mm] in diameter), then the thickness of the boundary layer around the hole orifices is approximately equal to the hydraulic diameter of the holes, so that viscosity cannot be neglected. This high viscous loss occurs as air passes through the perforations, leading to high absorption magnitude. The acoustic impedance model of MPP, separated by distances much larger than their diameters (radius), but small compared to the wavelength of impinging sound wave, was derived as [38, 42],

$$Z_{\text{MPP}} = \frac{8\mu t}{\phi \rho c r^2} \left(\sqrt{1 + \frac{x^2}{32} + \frac{\sqrt{2} x r}{16 t}} \right) + j\omega \left[\frac{t}{\phi c} \left(1 + \frac{1}{\sqrt{9 + x^2/2}} + 1.7 \frac{r}{t} \right) \right], \quad (2.21)$$

where $x = r\sqrt{\omega\rho/\mu}$, $\forall x \in [1,10]$, μ is the dynamic viscosity, t is the panel thickness, and ϕ is the porosity of the panel, and r is the radius of the hole. Taken into consideration end corrections due to geometric discontinuities at the inlet and exit of the perforation, Eq. (2.21) becomes

$$Z_{MPP} = \frac{8\mu t}{\phi\rho c r^2} \left(\sqrt{1 + \frac{x^2}{32} + \frac{\sqrt{2} x r}{4 t}} \right) + j\omega \left[\frac{t}{\phi c} \left(1 + \frac{1}{\sqrt{9 + x^2/2}} + 1.7 \frac{r}{t} \right) \right]. \quad (2.22)$$

The total impedance of the MPP can thus be defined as the sum of the impedance of the panel, and that of the air cavity behind it. i.e., $Z = Z_{MPP} + Z_{cavity}$. For each perforation, the cavity acts as a long tube driven by harmonic plane wave source (e.g., a piston) at $x=0$, and terminated by rigid boundary at $x=d$, as shown in Fig. 2.4.




Figure 2.4: An illustration of a single perforation in an MPP where each perforation serves as the neck of a Helmholtz resonator, and a long tube acting as the cavity which is backed by a rigid boundary at $x=d$.

The acoustic impedance of the cavity is given by

$$Z_{cavity} = -j\rho_0 c \cot(kd). \quad (2.23)$$

2.3 Metamaterials: An overview

The foundation of metamaterials research originates in the electromagnetic domain, particularly with the theoretical work of Veselago [43], who first discussed materials characterised by simultaneously negative permittivity and permeability. Such hypothetical materials were predicted to exhibit negative refraction, paving the way for future research on artificial media with tailored properties. At the time, these kind of materials were only a theoretical construct/idealisation with no special name given to them. R.M Walser [44] first coined the term for these materials as 'metamaterials' at the University of Texas, Austin, and defined it as *macroscopic composites having a man-made, three dimensional, periodic cellular architec-*

ture designed to produce an optimised combination, not available in nature, of two or more responses to specific excitation.

In 1995, Pendry et al.[45] constructed a novel material from a periodic array of thin metallic wire as the building microstructure, and were able to demonstrate that the new structure could depress plasma frequency in the GHz band, which is not achievable with any existing material. Following this theoretical demonstration, this concept was experimentally realised in which negative permeability material was obtained from copper etched ‘split-ring’ type resonators [46]. Shelby et al.[1] experimentally realised negative values of refractive index over a range of frequencies using a periodic array of copper split-ring resonators and wire strips on a fiberglass circuit board as shown in Figure 2.5, and called their fabricated structure *left-handed* metamaterial. Refraction is the principle upon which lenses and imaging operate, and according to Snell’s law, a ray crossing the boundary of two media must be refracted (slightly bent towards the normal to the surface, but not on the same side as the incident ray) at some angle. According to Snell’s law, all naturally occurring materials would possess a positive refractive index, however, Shelby’s experiment proved the possibility to engineer a material with unusual property of negative refraction (at certain frequencies) with the refracted ray on the same side as the incident ray. The two critical parameters of a material that are used



Figure 2.5: A periodic array of copper split-ring resonators on fibre glass circuit board [1].

to control electromagnetic waves in an electromagnetic medium are the permittivity, ϵ , and permeability, μ . The refractive index, n , is thus defined as in Eq. (2.24).

$$n = \sqrt{\mu\epsilon}. \quad (2.24)$$

A negative index material is realised when both ε and μ are simultaneously negative. Following these research works, the field of metamaterial has expanded and found applications in other fields such as in optics [47, 48] and acoustics [2, 49, 50], and it took about three decades for the first AMM to be fabricated and experimentally validated.

2.3.1 Photonic crystals

Yablonovitch [51] proposed the idea of a "photonic bandgap" in a periodic dielectric structure, drawing parallels to electronic bandgaps in semiconductors. This concept laid the foundation for photonic crystals. Yablonovitch and Sajeev [52, 53] then independently introduced the concept of photonic bandgap materials and proposed that periodic dielectric structures could prohibit the propagation of certain wavelengths of light, leading to bandgaps in the optical spectrum.

Researchers began conducting experiments to verify the existence of photonic bandgaps. The first experimental demonstration of a two-dimensional photonic crystal was achieved by Yablonovitch and colleagues [54] to validate the predict bandgap behavior. Kurt Busch and Sajeev [55] later extended the concept of photonic crystals to three dimensions. They introduced the idea of three-dimensional periodic structures that could exhibit complete photonic bandgaps, opening up new avenues for light manipulation. The development of fabrication techniques at the time played a crucial role in advancing photonic crystals. Electron beam lithography and other nanofabrication techniques enabled the precise structuring of periodic dielectric materials at the submicron scale. Throughout the late 1990s and early 2000s, research in photonic crystals expanded rapidly, with experimental efforts focused on improving fabrication techniques, achieving higher-quality structures, and exploring novel materials.

In optics, engineered structures designed to manipulate electromagnetic waves are called photonic crystals, and were first presented in 1987 [51, 53] for controlling optical waves in refraction index-modulated periodic structures. Photonic crystals are periodic structures that control the propagation of electromagnetic waves, particularly in the optical or near-infrared frequency range. These crystals are engineered to exhibit specific photonic bandgaps, which are frequency ranges where the transmission of light is strongly prohibited. Photonic crystals rely on the periodic arrangement of materials with contrasting refractive indices to achieve these bandgaps. The design of photonic crystals to achieve unusual wave manipulation requires tailoring the arrangement and properties of the constituent materials in order to realise negative values of effective permittivity and permeability. The resulting materials are

so-called *left-handed or negative-index* photonic crystals, and has led to applications including negative refractions [56], inverse Doppler effects [57] and superlens [58].

In the 2000s, the field of metamaterials emerged, encompassing both photonic and electromagnetic metamaterials. Metamaterials extended the concept of photonic crystals by introducing subwavelength unit cells and achieving unconventional optical properties.

2.3.2 Phononic crystals

Phononic crystals emerged in the 1990s, building upon the theoretical foundations of electronic band structures and photonic crystals, where researchers began exploring the possibility of manipulating the propagation of elastic waves, such as sound and vibrations, in periodic structures.

In the early 1990s, researchers like Sigalas et al. [59, 60] delved into the study of elastic and acoustic wave band structures in periodic systems. They developed mathematical models and theoretical frameworks to analyze the behavior of elastic waves in two-dimensional and three-dimensional structures. Kushwaha et al.[61, 62] introduced the concept of phononic bandgaps, they showed that by using a periodic arrangement of materials with different acoustic properties, it is possible to prohibit the propagation of elastic waves of specific frequencies, resulting in bandgap formation. Following these theoretical developments, the first phononic bandgap was experimentally demonstrated in 1998 by Del [63] , Sanchez [64] and Robertson[65], and since then, a plethora of work has been done and researchers began fabricating and characterizing phononic crystal structures using various materials and techniques [66–68].

Phononic crystals are periodic structures used to control acoustic waves/create acoustic bandgaps. The periodic arrangement of materials in a phononic crystal creates a band structure that governs the propagation of mechanical waves. This band structure consists of allowed and forbidden frequency ranges known as bandgaps. However, the presence of scatterers can be used to modify the band structure and hence alter the behavior of waves within the crystal. These scatterers are specific elements or features within the crystal structure that cause scattering of mechanical waves. Scatterers can take various forms, such as impurities, dislocations, defects, or engineered inclusions. These features typically have different acoustic properties compared to the surrounding materials, leading to the scattering of waves incident upon them. Phononic crystals are characterised by elastic bulk moduli and mass densities [61, 69]. The propagation of acoustic waves in a phononic crystal is governed by

the Bloch or Floquet theorem [70, 71] so that, when an acoustic wave propagates through a periodic structure, constructive interference leads to band gaps (stop bands) that inhibit wave propagation through the structure at the target frequencies. In this state, all the incident energy of the wave gets reflected in the source. In the case of destructive interference, all the incident energy is transmitted through the structure, resulting in the formation of passbands that allow wave propagation. Band gaps of specific target frequencies can be realised by tuning the cell's structural parameters. The width of the band gap increases with increase in the difference in mass density between the material of the scatterers and that of the host matrix [72]. Thus, the mass density, elastic modulus, shape, location, and topology of the scatterers, are essential parameters for manipulating acoustic wave propagation in the structure [73], and depending on their size, shape, and arrangement, scatterers can influence the transmission, reflection, and localization of mechanical waves. Some notable applications of phononic crystals include, acoustic and vibration control, waveguiding [74–76] and filtering [77–79], negative refraction [80, 81], sound focussing [82], energy harvesting [83–85] and acoustic sensing [86, 87].

With the passage of time, advancements in theoretical modeling, computational techniques, and fabrication methods (such as the 3D printing technology) have allowed for more complex and sophisticated phononic crystal designs. Researchers explored three-dimensional phononic crystals, investigated the existence of full bandgaps in various configurations, and studied the effects of defects, interfaces, and nonlinearity on phononic crystal behavior. More so, the development of AMMs with sub-wavelength units, whose underlying mechanisms enable unusual properties based on the interaction between the incident sound waves and the localized resonators within the metamaterial, leading to unprecedented material behaviour and acoustic properties, are currently receiving intense interest and attention among researchers.

2.3.3 Acoustic metamaterial

AMMs represent a rapidly evolving field of research that draws inspiration from the broader class of electromagnetic metamaterials (EMMs). Defined broadly, AMMs are artificially engineered structures designed to manipulate the propagation of acoustic waves in ways that surpass the capabilities of conventional materials. These metamaterials often rely on subwavelength-scale structural features that result in unusual effective acoustic parameters, such as negative mass density, negative bulk modulus, or both simultaneously. Since their

inception, AMMs have enabled novel acoustic phenomena including negative refraction, sub-wavelength imaging, acoustic cloaking, and advanced sound absorption.

The invention of metamaterials generally originated from the desire to overcome the limitations of conventional materials, explore unconventional behaviors, tailor the material properties, and unlock new possibilities in applications and technologies. While metamaterials originated from the field of electromagnetics, AMM is inspired by its electromagnetic counterpart to manipulate acoustic wave. AMMs have thus evolved with massive interest among researchers and the industry in recent years, and this has significantly broadened its scope and application areas. These materials possess the bulk acoustic properties of conventional materials on a cellular level, but behave as a homogeneous system on a macroscopic level, having a unique effective bulk properties not present in its constituent building units [88].

The term 'metamaterial' comes from the fusion of two words: 'meta' and 'material'. In ancient Greek, 'meta' means beyond, and 'material' means matter. Hence, beyond conventional material or matter. That said, metamaterial can be defined as any material that is engineered to have a property that is uncommon in conventional materials. Thus, due to the limitations of conventional acoustic materials, new materials have emerged which can control and manipulate sound waves like no other existing materials. AMMs are artificial materials composed of periodic arrangement of sub-wavelength units engineered to manipulate sound in an unconventional manner. The unit cells constitute the basic building blocks which are designed to have the desired properties. The extra-ordinary sound wave manipulation are achieved by controlling the two critical bulk acoustic properties of a material: adiabatic bulk modulus, B , and effective mass density, ρ . This is based on the linear acoustic wave equation which governs the propagation of sound wave in a homogeneous fluid medium as in Eq. (2.25)

$$\nabla^2 p - \frac{1}{c^2} \frac{\partial^2 p}{\partial t^2} = 0, \quad (2.25)$$

where p is acoustic pressure, $c = \sqrt{\frac{B}{\rho_o}}$ is the thermodynamic speed of sound, and ρ_o is the mean density of the medium. From Eq. (2.25), the overall behavior of the propagating wave is entirely dependent on ρ and B of the material. $B = \rho_o \left(\frac{\partial P}{\partial \rho} \right)_{\rho_o}$, measures the resistance to compression or expansion of the fluid medium, while ρ or the effective mass density can be considered as the ratio of force acting on a body and the net acceleration per unit volume produced in the body, in the direction of the force, based on Newton's second law. Thus,

to achieve extra-ordinary material property, either or both of B and ρ must be negative, analogous to those of electromagnetic material where the permittivity and/ permeability of the material must be negative. Thus, an AMM can be of the type negative B or negative ρ or negative B and ρ simultaneously. In the regions of negative ρ and/or negative B , materials violate the mass-frequency law and remarkable sound manipulation is achieved. It is noteworthy that the unit cells of AMMs do have conventional or positive values of ρ and B , but the geometry and the way in which these cells are arranged within the lattice results in an unconventional homogeneous material property on a macroscopic scale. Similar to the concept of negative index materials that were used to control electromagnetic waves, a similar concept was borrowed to control acoustic waves.

Liu et al. [2] fabricated the first AMM in the year 2000, and was called sonic crystals. This type of AMM had negative B values. The unit cell was a 1 [cm] diameter hard lead ball coated with a 2.5 [mm] layer of silicone rubber as shown in Figure 2.6. It worked on the principle of negative bulk modulus, B . The unit cells were arranged in an $8 \times 8 \times 8$ cube as shown on the right figure in Fig. 2.6. When sound wave impinged on the cubic structure, a transmission coefficient curve, τ , generated from the wave-structure interaction is produced as shown in Fig. 2.7. The individual balls were found to obey the mass-frequency law with conventional bulk acoustic properties, however, when the balls were arranged as shown in Fig. 2.6, it resulted in the curve in Fig. 2.7 which violates the mass-frequency law especially at around 400 [Hz] and 1100 [Hz], where we have sudden drops in transmission coefficient. This means that the material can provide exceptional performance with high transmission loss at those frequencies.

The second AMM was proposed in 2004 by Liu and Chan [89]. It was based on simultaneous negative B and ρ . The idea here was similar to Veselago's medium in electromagnetism and exhibited various interesting properties, including a negative refractive index. In acoustics, the double negativity arises from low-frequency resonances, similar to that of electromagnetism. However, unlike electromagnetism where the negative permeability and negative permittivity stem from different resonance mechanisms, in acoustics, the negative density and modulus are derived from a single resonance structure. Although the concept of double negative acoustic material was theoretically proven, no experimental validation was provided.




Figure 2.6: Sonic crystal: First AMM fabricated [2].




Figure 2.7: A transmission coefficient curve generated from the first $8 \times 8 \times 8$ sonic crystal [2].

In 2008 [3], the first complete model of a membrane type AMM came into existence. The unit cell is a stretched membrane with a centre mass clamped within a subwavelength wave guide as shown in Fig. 2.8. It works on the principle of negative mass density. For a 6 [mm] diameter circular steel disk of 300 [mg] fixed at the centre of the stretched membrane made of rubber, the observed transmission coefficient curve (on the right plot of Fig. 2.8 expressed in percentage, shows a sharp drop in transmission coefficient at extremely low frequency, which is contrary to what would be expected from a conventional acoustic material.




Figure 2.8: Membrane type AMM [3].

2.4 Acoustic metasurface

Over the years, there has been growing interest across various technological fields in developing smaller, lighter, and more compact devices and system components. Miniaturisation is evident in sectors such as consumer electronics, healthcare, and aerospace, where it has revolutionised industries and opened up new possibilities. The demand for compact systems is driven by the need for portability, enhanced performance, energy efficiency, and the integration of multiple functionalities into smaller form factors. In line with this trend, this research adopts the metasurface concept to enable design simplicity, lower fabrication complexity, and compact device footprints.

The field of metasurfaces is driven by the need to control sound wave propagation using miniaturised, compact devices. Compared to their bulk AMM counterparts, metasurfaces offer added value and enable unconventional functionalities [90]. Acoustic metasurfaces are essentially 2D realisations of bulk AMMs, designed to control and manipulate sound with high precision. While traditional metamaterials are typically 3D structures, metasurfaces operate in a single plane, allowing for more compact and surface-conformal implementations.

The class of thin 2D metasurfaces was introduced since 2011 for optical waves via the generalised Snell's law, and this has unearthed novel opportunities for optical wave manipulation and other applications [91], and since acoustic waves also follow the Snell's law, a plethora of research works [4, 5, 7, 92–107] have adapted this material for sound control, and also opened up new opportunities to different applications in the field of acoustics. However, major challenges emerge in translating the success of electromagnetic metasurface to the acoustic domain, and this lies in the mechanical nature of sound and the wide dissimilarity between the wavelength of electromagnetic waves and that of acoustic waves [7, 92]:

the longer wavelength of acoustic waves require a thicker structure for the wave to effectively interact with the structure and produce the desired acoustic behaviour.

Acoustic metasurface concept is based on the arrangement of sub-wavelength unit in a 2D plane, and several design concepts have been proposed based on, but not limited to, Helmholtz resonators [108], membranes [95] and coiling-up space structures [7, 96, 109]. With extraordinary sound absorption capabilities, the uniqueness of metasurfaces lies in their ability to freely steer the wave fields such that the phase and/or amplitude is utterly controlled, and one vital reason for the interest in acoustic metasurfaces is to overcome the challenge of using bulk AMMs to control acoustic wave of long wavelengths in air and water [90]. Following from the success of bulk AMMs [110–113], acoustic metasurfaces have significantly advanced the field of acoustic wave manipulation, enabling the design of compact lightweight structures with unprecedented functionalities. Thus, use of metasurfaces with optimised geometries can well overcome the challenge of bulk AMM with lightweight compact structures

Owing to the millimetre/submillimeter sizes of acoustic metasurfaces' sub-wavelength units, one obvious challenge would be the fabrication of these sizes using traditional fabrication methods. However, with the rapid advancements in additive manufacturing, these sub-wavelength units can be easily manufactured to sufficient precision, enhancing straightforward experimental validations of fabricated structures. And with the high air-solid impedance mismatch, polylactic acid (PLAs) or resin based 3D printers (such as fused deposition modelling (FDM) or Stereolithography) commonly used are able to provide a highly reflective boundaries for the sound wave.

2.4.1 Physics of acoustic metasurface

Acoustic waves incident on a planar boundary between two media obey the Snell's law, which describes the relationship between the angles of incidence and the angle of refraction when a wave crosses the boundary between two different media, and these two angles are dependent on the difference in refractive indices between the two media. Of course this law is based on the understanding that waves are continuous and coherent phenomena, and their behavior is governed by the principles of wave propagation. In other words, the change in phase of a wave as it crosses the boundary between two media occurs smoothly and continuously as the wave propagates. This assumption allows for the mathematical description of the relationship between the angles of incidence and refraction in Snell's law to be possible. Snell's law can thus be leveraged to control the reflected and transmitted acoustic waves beyond what is

obtainable with natural medium boundaries. This extraordinary control of acoustic waves is realised by introducing instantaneous phase shifts in the acoustic propagation path by arranging periodic sub-wavelength structures on a planar interface [90]. By carefully designing the geometry, size, and arrangement of the sub-wavelength structures, acoustic meta-surfaces can achieve diverse functionalities such as abnormal reflection, refraction, absorption, wavefront modulation, etc.




Figure 2.9: Schematic illustration of generalised Snell's law and acoustic metasurface. The interface is treated with an engineered periodic structure which results in an instantaneous phase shift along the acoustic path, and the generalised Snell's law of reflection and refraction occurs at the interface

Consider an acoustic plane wave impinging on the interface between two media treated with a periodic material such that an additional phase shift, which is a continuous function, $\Phi(x)$, of the position along the interface, is introduced to the acoustic wave which is then transmitted or reflected from this interface as shown in Fig. 2.9. If the interface lies on the x-y plane ($z=0$), then a phase shift, $\Phi(x)$, for a 1D interface can be created for the incident wave as it is reflected from or passes through the interface. Using Fermat's principle which states that the acoustic path length (product of the physical length and the refractive index of the medium) of an acoustic ray travelling from point A to point B after reflection from the interface (or to point C after refracting into a second medium) is a minimum (see Fig. 2.9). Following wave theory, the difference in acoustic path length is equivalent to the variation of the propagation phase and according to Fermat's principle, the path traversed by the acoustic wave between two points A and B for instance, should be the path of least phase change.

When an acoustic ray travels from point A at coordinates (x_A, z_A) and impinges the interface at point O with coordinates $(x, 0)$ at angle of incidence, θ_i , and then reflects to point B at coordinates (x_B, z_B) , the overall phase change of the acoustic wave along this path is given by

$$\Phi_r(x) = \phi(x) + k_1 \sqrt{(x - x_A)^2 + z_A^2} + k_1 \sqrt{(x_B - x)^2 + z_B^2}, \quad (2.26)$$

where $k_1 (= \frac{2\pi}{\lambda_1})$ is the wavenumber in medium 1, and the instantaneous phase change, $\phi(x)$, appears in Eq. (2.26) due to the contribution from the engineered interface. Differentiating $\Phi_r(x)$ with respect to x and equating the resulting expression to zero for minimum acoustic path length yields

$$\frac{d\Phi_r(x)}{dx} = \frac{d\phi(x)}{dx} + \frac{2\pi(x - x_A)}{\lambda_1 \sqrt{(x - x_A)^2 + z_A^2}} - \frac{2\pi(x_B - x)}{\lambda_1 \sqrt{(x_B - x)^2 + z_B^2}} = 0, \quad (2.27)$$

where $\lambda_1 = \frac{c_1}{f}$ is the wavelength of sound in medium 1, and c_1 and f are the speed of sound and the frequency, respectively. Upon simplification, Eq. (2.27) becomes

$$\frac{d\phi(x)}{dx} + \frac{2\pi}{\lambda_1} (\sin\theta_i - \sin\theta_r) = 0. \quad (2.28)$$

where, based on Fig. 2.9, $\sin\theta_i = \frac{(x - x_A)}{\sqrt{(x - x_A)^2 + z_A^2}}$ and $\sin\theta_r = \frac{(x_B - x)}{\sqrt{(x_B - x)^2 + z_B^2}}$. Rearranging Eq. (2.28) yields Eq. (2.29).

$$\sin\theta_r - \sin\theta_i = \frac{\lambda_1}{2\pi} \frac{d\phi(x)}{dx} \quad (2.29)$$

The reflection of acoustic waves off the interface is governed by the generalized Snell's law of reflection as in Eq. (2.29), where a nonlinear relationship exists between θ_i and θ_r . It is noteworthy that Eq. (2.29) contradicts the conventional law of reflection which requires that $\sin\theta_i - \sin\theta_r = 0$. In the same vein, applying the Euclidean distance between two coordinate points from Figure 2.9, the total phase change for the transmitted wave from point A to point C through refraction at the interface point O is given by

$$\Phi_t(x) = \phi(x) + k_1 \sqrt{(x - x_A)^2 + z_A^2} + k_2 \sqrt{(x_C - x)^2 + z_C^2}, \quad (2.30)$$

where the coordinates of point C is (x_C, z_C) , $k_2 (= \frac{2\pi}{\lambda_2})$ is the wave number in medium 2, and $\lambda_2 (= \frac{c_2}{f})$ is the wavelength of the sound in medium 2. The relationship between θ_i and

θ_t can be obtained by differentiating Eq. (2.30) and equating the resulting expression to zero for the shortest acoustic path from A in medium 1 to C in medium 2.

$$\frac{d\Phi_t(x)}{dx} = \frac{d\phi(x)}{dx} + \frac{2\pi(x - x_A)}{\lambda_1 \sqrt{(x - x_A)^2 + z_A^2}} - \frac{2\pi(x_C - x)}{\lambda_2 \sqrt{(x_C - x)^2 + z_C^2}} = 0, \quad (2.31)$$

Simplifying Eq. (2.31) and rearranging gives the generalised Snell's law of refraction

$$\frac{1}{\lambda_2} \sin \theta_t - \frac{1}{\lambda_1} \sin \theta_i = \frac{1}{2\pi} \frac{d\phi(x)}{dx}. \quad (2.32)$$

If $\phi(x)$ is independent of x , then Eq. (2.29) and Eq. (2.32) follow the conventional Snell's law, i.e., $\frac{d\phi(x)}{dx} = 0$. Now suppose that $\phi(x)$ is a linear function of x such that $\phi(x) = bx$, then Eq. (2.29) and Eq. (2.32) become, respectively

$$\sin \theta_r - \sin \theta_i = \frac{\lambda_1}{2\pi} b. \quad (2.33)$$

$$\frac{1}{\lambda_2} \sin \theta_t - \frac{1}{\lambda_1} \sin \theta_i = \frac{1}{2\pi} b. \quad (2.34)$$

Here, the reflected wave has a specific value θ_r , and the transmitted wave also has a specific value θ_t , resulting in an anomalous reflection and refraction governed by the generalised Snell's law. Thus, θ_r and θ_t can be controlled using the parameter, b . Also of note is the unconventional phenomena that $+\theta_i$ and $-\theta_i$ do not give the same θ_r and θ_t according to the generalised Snell's law. In the same vein, if $\phi(x)$ is a non-linear function of x such that, say $\phi(x) = \alpha x^2$, then Eq. (2.29) and Eq. (2.32) become, respectively

$$\sin \theta_r - \sin \theta_i = \frac{\lambda_1}{\pi} \alpha x \quad (2.35)$$

$$\frac{1}{\lambda_2} \sin \theta_t - \frac{1}{\lambda_1} \sin \theta_i = \frac{1}{\pi} \alpha x \quad (2.36)$$

Here, the reflected and refracted angles depend on the position, x , along the interface. This provides the opportunity to manipulate the reflected and transmitted acoustic beams by tuning the nonlinear function such that the local acoustic response at each point behaves differently [90].

In practice, planar structures that are capable of introducing instantaneous phase shift in the path of wave propagation can be engineered from a single layer of unit cells with

periodically arranged units on a 2D plane. Such structures are called metasurfaces. The periodic arrangement of these subwavelength units on a planar interface is vital to realising the desired acoustic behaviour. The periodicity determines the interaction of the sound waves with the metasurface. Furthermore, it is noteworthy that acoustic metasurfaces do not necessarily have near-zero thicknesses but rather sub-wavelength thicknesses, and the sub-wavelength units constituting the metasurface can be of any shape, and can only have vanishing thicknesses when compared to the incident wavelength [90].

While the concepts and principles from optics provide the foundation for acoustic metasurface, directly applying these concepts to acoustic waves is not straightforward due to inherent differences between the two. Thus, developing and implementing acoustic metasurfaces require specific adaptations and considerations to suit the properties and behavior of sound waves. The fundamental differences between electromagnetic (optical) waves and acoustic waves that need to be taken into account include:

- **Nature:** Electromagnetic waves are transverse waves, meaning that particle oscillations occur in a direction perpendicular to the direction of wave propagation. Acoustic waves, on the other hand, are longitudinal waves, with particles oscillating in the direction of wave propagation.
- **Speed and Wavelength:** Electromagnetic waves, such as light, typically have much higher speeds and shorter wavelengths compared to acoustic waves. The wavelength of light is typically on the order of hundreds of nanometers, while sound waves have wavelengths on the order of meters or larger.
- **Interaction:** Electromagnetic waves consist of oscillating electric and magnetic fields that can interact with matter primarily at the atomic and molecular level, involving the interaction of the electromagnetic fields with charged particles and the electronic structure of matter. Acoustic wave on the other hand are mechanical waves that require a material medium to propagate. They interact with matter primarily through the compression and rarefaction of the particles of the fluid medium, leading to effects such as absorption, reflection, and refraction.

Furthermore, there are acoustic phenomena which are not present in optical systems. For example, the frequency-dependent thermo-viscous losses are not uncommon in acoustic systems with narrow channels, and their impact on the performance of acoustic metasurfaces

has been investigated [114]. Another acoustic effect is the fluid–structure interaction at the wall (or non-rigid boundary) of an acoustic metasurface [115], although this has been said to have a less significant impact than thermo-viscous losses. That said, research works in [8, 116, 117] have shown that acoustic metasurfaces may still be a challenging task to realise, and thus requires the development of a fundamentally new physics beyond the generalised Snell’s law upon which optical metasurface was developed.

2.5 Review of acoustic metasurface designs

A number of acoustic metasurface design concepts have been explored in the literature, including the coiling-up space structures [4, 93], the hybrid structure concept comprising Helmholtz resonators and straight pipes [108], membrane-type metasurfaces with hybrid resonances [95]. Designs of acoustic metasurfaces focused on controlling one of reflected, transmitted or absorbed waves, require a precise control of the phase and amplitude of the wave propagation, and this can be realised by exploiting existing AMM structures including Helmholtz-resonator-like structures, coiling-up space structures and membrane-type structures [90].

2.5.1 Acoustic metasurfaces for controlled reflection

The coiling-up space metasurfaces with subwavelength dimensions creates a longer propagating path for the acoustic wave, such that the wave travels through a longer path beyond the physical dimension of the structure. A simple space coiling labyrinthine structure can achieve extreme elongation of the propagation path by coiling-up space [118–123]. The longer path length, which is achieved by introducing additional spatial loops or twists in the metasurface structure, allows to modulate or control the properties of the reflected wave such as its amplitude, phase, or direction, and also allows to adjust the reflected wave to have any phase shift in the range $[0, 2\pi]$, offering a wide range of possibilities for controlling the sound wave.

A coiling-up space structure whose propagation path can be tuned is a viable approach to design a metasurface of the acoustic-reflection type [4, 93, 96, 124–127]. The basic labyrinthine unit as shown in Fig. 2.11(a), forms the fundamental building unit of an acoustic metasurface for controlling sound reflection. The acoustic wave enters the the labyrinthine structure through the inlet, propagates through the coiled channel and is completely reflected at the hard end, traversing the same path back to exit at the inlet. This constitutes the reflected

wave. In a free space, the path length of the propagating wave would be the sum of the straight distances of the forward and backward propagation, i.e., $2a_x$, as shown in Fig. 2.11(a). However, with the introduction of the coiling-up space structure into the medium, the wave zigzags through the structure, and its propagating length is significantly increased. Hence, the phase of the reflected wave due to the labyrinthine structure is significantly delayed, which enables full control of the phase shift. And by varying the geometric parameters of the structure, the phase shift can be controlled over a full range [90]. In addition to the labyrinthine structure, a planar surface realised from Helmholtz resonators can generate double reflections, namely, an ordinary reflection and an extraordinary one that obeys the generalized Snell's law of reflection [125].

Acoustic metasurfaces with compact/lightweight geometries are expected to manifest new functionalities through reflected wavefront manipulation. One such planar design is the ultra-thin Schroeder diffuser proposed in [5], which is based on the acoustic metasurface concept, and is composed of a 2D array of locally resonant Helmholtz-like structures each carefully designed to produce a specific phase shift according to a certain quadratic residue sequence delineated by Schroeder's theory. The unit cell of this system is a locally resonant element, and its acoustical response can be tuned flexibly and precisely by adjusting a single geometrical parameter. The design focus of the system is to have minimal thickness while ensuring that the performance is not affected significantly by thermo-viscous losses.

2.5.2 Acoustic metasurface for controlled transmission

The problem of impedance matching is very key when designing acoustic metasurfaces for manipulating transmissive acoustic waves, which may not be considered in the design of reflective metasurfaces [90]. Acoustic metasurfaces of the transmissive type are able to manipulate the transmitted wavefront curvature across the metasurface, and provide a precise control over the transmission properties of sound waves such as the amplitude, phase and direction. Solids generally do not allow the propagation of sound through them due to the high impedance mismatch between the ambient air and the solid and this reduces the transmission efficiency of the acoustic system. That said, the design of a transmissive acoustic metasurface are tailored towards minimising the impedance mismatch between the metasurface and the ambient air.

Unconventional refraction of acoustic waves at the interface between two media demonstrates how a transmissive acoustic metasurfaces may be realised (see Fig. 2.9). Phase change




Figure 2.10: Acoustic metasurfaces based on controlling the reflected wave by wavefront manipulation. Figure (a) is the phase shifts of the reflected waves from four different types of labyrinthine units in Figure (b), where Figure (b) is a 3D printed labyrinthine structure with eight labyrinthine units. The label (m, n) indicates that m plates from the top and n plates from the bottom are used for a labyrinthine unit. The eight dots in Figure (a) refer to the chosen eight units shown in **B** [4]. Figure (d) is a schematic of the Schroeder diffuser. The top and bottom images are the top and 45° angle views of, while the insets is the cross section of unit cells in SD and MSD, with the thicknesses being $\frac{\lambda_0}{20}$ [5]




Figure 2.11: Metasurface structures of the (a) coiled-up type, (b) Helmholtz type, and (c) membrane type

across the interface is gradual and continuous according to the classical Snell's law, however, the additional phase gradient introduced by acoustic metasurfaces create an anomalous refractive angle, different from the classical angle as shown in the generalised Snell's law of Eq. (2.32). This theory shows that an engineered surface can achieve various forms of wavefront modulation, such as wave focusing [119, 128, 129] and wave deflections [6, 7, 94], etc.

The design of an acoustic metasurface unit composed of four Helmholtz resonators to control the phase, and a straight pipe to induce hybrid resonances for impedance mismatch compensation, was the first realisation of a high refractive index and high transmission structure (Fig. 2.11(b)). This metasurface possess an adjustable phase velocity and a near-perfect transmission efficiency, and features a deep sub-wavelength width. Since then, a number of designs have been proposed to realise transmissive acoustic metasurfaces, such as the tapered labyrinthine metamaterials [130], spatially varied coiling-slit units [6, 131] and impedance-matching layers of gradient index [132].




Figure 2.12: Fabricated supercell of an acoustic metasurface structure (top image) for anomalous deflection of transmitted acoustic wave [6]. Composed of spatially varied coiling-slit subunits, with a deep subwavelength thickness. By optimizing its microstructures, high performance wavefront steering is exhibited. (b)-(d) Measured amplitude, temporal, and phase field distributions. The green arrows and parallel lines indicate the propagating wavefronts

Researchers [7, 130] proposed a horn-like labyrinthine transmissive metasurface with a spiral-shaped geometry and a high transmission efficiency. The fabricated structure is a supercell containing six types of such horn-like labyrinthine cells with different levels of spirallings as shown in Fig. 2.13 is able to span a phase change range of 2π (i.e., $[-\pi, \pi]$) while maintaining a uniform sub-wavelength thickness in the frequency range of 2500–3500 [Hz].




Figure 2.13: Fabricated supercell of a tapered labyrinthine metasurface structure (left) and phase changes corresponding to six types of unit cells [7].

2.5.3 Acoustic metasurfaces for controlled absorption

Sound absorption for noise reduction is an important application of AMM, and has received massive attention among researches and the industry, owing to the detrimental effect of noise to humans, animals, and also to machines and system components [88]. Currently, transportation generally contribute the bulk of the noise problem, with low-frequency noise being especially pernicious and challenging for traditional acoustic treatment. Traditional sound-absorbing materials have fixed absorption spectra, and are tuned by manipulating their thicknesses which can become very bulky and not necessarily effective at low frequency [133]. It is therefore imperative to develop an advanced class of sound absorbers that are compact, lightweight, and subwavelength in thickness, and that can be embedded within structures and reconfigured as needed.

Sound absorption in the low-frequency range ≤ 500 [Hz] have been a scientific and technological challenge for engineers. To be specific, the absorption of low-frequency noise is a difficult one because of the weak dissipation of conventional materials in the low-frequency regime [134], which include visco-elastic materials, porous and fibrous materials [135], gradient-index materials [91, 136] and micro-perforated plates with backing cavities [41].

A surface can be engineered to eliminate reflections by using, for instance, the configuration in Fig. 2.11(c). The resulting metasurface realises perfect absorption at a frequency close to the anti-resonance despite the deep-subwavelength thickness, using hybrid resonance concept [90]. It is noteworthy that several designs have achieved narrow band sound absorptions with thin absorbers without any complimentary dissipative mechanism [41, 137], these designs dissipate acoustic energy using the thermo-viscous effect occurring in the nar-

row channels of the system. However, the membrane-based structures enable large change in membrane vibration, resulting in a full absorption when a very weak absorption coefficient is amplified by a large energy density [90].

The field of AMM has received a surging interest and development since its inception in 2000, and has grown to include a range of functionalities and novel applications, such as sound absorption. There are a few demonstrations of sound absorption for low-frequency application [112, 138–141] which are mostly based on two configurations: membrane-type and plate-type AMM. However, these types are either weak or operate a little above the targeted frequency range [90].

Sound-absorbing structures designed to have specific properties, such as using sound dissipating materials, resonant chambers, or employing geometrical configurations that dampens sound energy, have been reported. Specifically, it discussed the relationship between absorption and the thickness of the structure and the approach to realising broadband acoustic absorbers [133]. These structures are designed to minimise sound reflections and to maximise absorption. Planar structures, on the other hand with deep subwavelength thicknesses can provide exceptional functionalities that have enabled the development of modern acoustic absorbers.

Acoustic metasurface for impedance matching

Acoustic metasurfaces with impedance-matched capabilities and hybrid resonances are capable of realising perfect absorption to incident sound wave. Such surfaces minimise reflection and transmission [95]. Such simple design may be based on Figure 2.11c, which incorporates a reflecting membrane resonator composed of an elastic membrane with a platelet and a reflecting hard surface. A gas with a small adiabatic index is sealed between the platelet and the hard surface. The decorated membrane resonator has two resonant modes: one related to the platelet's oscillation and another related to the surrounding elastic membrane's oscillation. The sealed gas and reflecting surface introduce additional impedance in series with the decorated membrane resonator, altering the resonance conditions. Consequently, the resonances of the decorated membrane resonator are compelled to combine, creating new hybrid modes with absorption conditions. In other words,, its working mechanism is that the amount of displacement in the membrane is separated from the radiation mode. Although the material has a low dissipation, the displacement in the membrane can be quite significant. This significant displacement results in full absorption when a small absorption

coefficient is multiplied by a high energy density. This absorbing metasurface is designed to match the impedance and is usually very thin, smaller than the incident wavelength. It can work effectively at very low frequencies, around 150Hz [95]. However, it is noteworthy that in certain situations, using these metasurfaces can be difficult due to their fragility and limited effectiveness across a narrow range of frequencies [90].

Acoustic metasurface with ultra-thin geometry

Extremely thin and highly efficient acoustic metasurface absorbers based on coiled-up structures are increasingly gaining traction. These absorbers leverage high-refractive-index materials to effectively slow down acoustic waves [8] using coiled chambers such as spirals, labyrinthine structures, etc. These structures essentially create phase lags along the propagation path [118, 121]. The primary purpose of this is to reduce the size of large structures to a scale much smaller than the wavelength of the acoustic waves. Additionally, it enables significant dissipation of the incident acoustic energy, resulting in a nearly perfect absorption. This absorption can be achieved when specific conditions of resonance and impedance-matching are met. Hence, in recent years, there has been an increasing focus on investigating acoustic absorption using coiled metasurfaces. This approach shows potential in addressing the difficulties associated with absorbing low-frequency airborne sound.

A deep sub-wavelength metasurface composed of a perforated plate and a coiled coplanar air chamber (Figure 2.14) was proposed in [8]. The metasurface realises a deep subwavelength thickness down to a feature size of $\approx \lambda/223$. The system parameters which achieved full absorption at 125 Hz are $d = 3.3$ [mm], $t = 0.2$ [mm], $w = 12$ [mm], $b = 1$ [mm] and $a = 100$ [mm]. In order to achieve perfect or near-perfect absorption, the metastructure was designed to fulfill two essential conditions: a resonant state and impedance matching with the surrounding air. These conditions ensure that all the acoustic energy is transferred into the coiled chamber instead of being reflected, and eventually absorbed within the perforated hole due to thermo-viscous losses. As the acoustic waves travel through the channel, the overall propagation length of the wave increases, resulting in a reduced sound velocity and an elevated acoustic refractive index. This facilitates the dissipation of acoustic energy by converting it into heat.

Another design of the coily type [9] explores a multi-coiled concept based on a coplanar coiled chamber with labyrinthine passages, and an embedded aperture, to achieve a near-perfect absorption at extremely low frequency of 50 Hz, and featuring an ultra thin thickness




Figure 2.14: Acoustic metasurface composed of a centred perforated plate and a coiled air chamber (left) and the corresponding acoustic response (right) at a resonant frequency of 125 Hz [8].

down to $\lambda/527$. The embedded aperture adds an extra degree of freedom to the system, which helps to shift the resonance frequency towards extreme low-frequency without increasing the total thickness of the metasurface and the effective length of the channel. The physics behind this metasurface is based on a hybrid resonance mechanism of coiling-up space geometry due to the spiral chamber (with labyrinthine passages), and the resonance of Helmholtz-like cavities formed by the corners of the spiral. The system parameters for the presented result in Figure 2.15 are, $a = 100$ [mm], $d = 6.5$ [mm], $h_a = 9$ [mm], $h_c = 12$ [mm], $w = 12$ [mm], $t = g = 2$ [mm].

Some other designs have also exploited the coiled-up (spiral structures) space approach to realise perfect absorption at low frequencies (≤ 500 Hz) [142, 143].

In general, the approach to designing acoustic metasurfaces depends on the desired physical properties and the targeted applications, which can include reflection, transmission, and absorption. It is therefore important to evaluate the design choices before addressing specific applications to ensure effective and appropriate functionality [90]. Metasurfaces based on Helmholtz resonator-like elements are advantageous because they are easy to fabricate and are adjustable [90]. They are mostly used for controlling reflection. These metasurfaces rely on local resonance, which makes them thinner and simpler in design compared to coiling-up space structures. On the other hand, coiling-up space structures are better suited for applications where sound absorption is key, and transmission is not required due to impedance




Figure 2.15: A multi-coiled metasurface absorber with an embedded aperture (left) and the corresponding acoustic absorption response (right) which achieves a full absorption at a resonant frequency of 50 [Hz] [9].

mismatch [90]. They generate dissipation in intricate channels and slow down the effective phase velocity. However, this concept requires a trade-off between metasurface thickness and thermo-viscous dissipation caused by the coiled channels. The choice of materials and fabrication process also adds complexity to the design. In the case of membrane-type metasurfaces, they are very thin and highly efficient for absorbing sound at low frequencies. However, they have the drawback of being delicate and more difficult to fabricate [90].

2.6 Deep learning optimisation

Before the use of data-driven techniques, the design of AMMs primarily relied on physics-based modelling and traditional optimisation algorithms. These approaches typically combined analytical or numerical simulations (e.g., finite element method (FEM), boundary element method (BEM)) with optimisation heuristics to improve key acoustic performance metrics such as absorption coefficient, transmission loss, or bandgap width. Given the computational cost and scaling limitations of traditional optimisation methods, especially in high-dimensional, multi-objective cases, researchers have increasingly turned to DL as an alternative approach.

Unlike traditional materials, the performance of AMMs is determined not only by constituent materials but also by the intricate geometric structuring at sub-wavelength scales. As AMMs evolve toward multi-functionality, compactness, and tunability, the design space

expands correspondingly. Traditional physics-based approaches, such as parameter sweeps, gradient-based optimisation, or topology optimisation—have contributed significantly to early AMM research, but they are computationally expensive and do not scale well for inverse design [144] or real-time applications. One of the primary challenges in AMM design is inverse mapping: given a desired acoustic response (e.g., absorption spectrum, transmission loss, or phase modulation), how can one determine the geometry or structural parameters that produce it? Traditional solvers must iterate through thousands of simulations, often guided by heuristics or evolutionary algorithms, to approximate a feasible design. These methods, although powerful, suffer from long convergence times and may fall into local optima.

DL offers a compelling alternative by enabling data-driven surrogate models that are trained (calibrate the DL model parameters) to learn the underlying structure–property relationships from data. Once trained, such models can perform near-instantaneous predictions of acoustic responses or generate optimal geometries for target functionalities. For example, Wang et al. [145] demonstrated the use of deep generative model that are able to generate thermal metamaterial geometries for the control, manipulation and processing of heat flow with high efficiency. Similarly, Yan et al. [146] implemented a hybrid deep neural network-genetic algorithm (DNN-GA) to optimise a multi-functional acoustic absorber for enhanced absorption efficiency across a wide frequency range, and was reported to achieve a 4000-fold improvement in computational efficiency compared to traditional FEM.

Another major benefit of DL is its ability to generalise across parameter spaces. Whereas traditional methods must recompute for every new design, a well-trained deep neural network can interpolate and even extrapolate, offering robust predictions over a wide range of configurations. Trinh et al. [147] trained a surrogate neural network on data generated from a multi-scale acoustic model based on the Johnson–Champoux–Allard–Pride–Lafarge (JCAPL) framework to map between microstructure and acoustic response in a physics-informed way. Their approach enables robust model identification and validation under tight data constraints and supports efficient exploration of design parameter spaces. Furthermore, the use of hybrid strategies, combining physics-based modelling with machine learning, has enhanced the interpretability and reliability of these data-driven models [148].

2.6.1 Review of deep learning models in AMM design

DL has spurred a paradigm shift in AMM design, enabling both forward modeling (predicting performance) and inverse design (identifying geometry/parameters for desired behavior).

These data-driven approaches dramatically reduce design cycle time and simulation cost, moving from iterative physics-based workflows to intelligent, real-time generation of AMM structures.

Liu et al. [149] proposed a novel DL-based design framework to develop an ultra-broadband AMM absorber capable of achieving near-perfect sound absorption across a wide frequency range. Recognising the limitations of conventional Fabry–Pérot (FP) channel arrays, i.e., arrays of straight acoustic ducts that rely on standing-wave resonances and therefore require physically long channels to achieve low-frequency absorption, particularly the complexity of accurately modelling the coupled resonance effects across numerous non-uniform channels, the authors employed a data-driven approach to circumvent the need for computationally intensive parametric optimization. A dataset comprising approximately 20,000 FEM-simulated configurations of 10×10 FP micro-structures was generated, with each structure defined by unique channel-length distributions. Two convolutional neural networks (CNNs) were trained: one for forward modeling (predicting absorption spectra from geometry), and another for inverse design (predicting geometry from target spectra). By cascading the networks, the authors rapidly identified metamaterial designs exhibiting average absorption rates exceeding 97% across 860–8000 [Hz] in simulation. The paper reported a training time of 10 [min]. To validate their results, the optimised structure was fabricated using high-resolution 3D printing and tested in an impedance tube. The experimental measurements confirmed good agreement with numerical predictions, achieving near-total absorption from 860–3200 [Hz], limited only by the tube’s upper frequency range. This study demonstrates the integration of DL and physical simulation in the inverse design of AMMs, offering a scalable and efficient route to broadband absorbers without reliance on traditional iterative or gradient-based optimisation techniques. The transmission of vibrations and structure-borne noises through plate and shell structures commonly occurs through the propagation of elastic flexural waves. Phononic metamaterials, which possess frequency band gaps, have proven effective in blocking elastic waves within specific frequency ranges. However, designing these metamaterials traditionally involves a time-consuming trial-and-error process [150]. Liu, Chan, and Wu [150] developed a deep-learning-based inverse design framework for phononic-plate metamaterials aimed at attenuating structure-borne flexural wave noise in the critical 2–4 [kHz] auditory frequency band. Employing the Mindlin plate theory for forward modeling, they generated only 360 FEM-based training samples, each defined by five

geometric slot parameters of a periodic unit cell. A deep neural network (DNN) was trained to predict unit cell geometry for a desired bandgap, achieving 2% error in center-frequency matching during testing. The optimised design exhibited ≈ -1 [dB/mm] omni-directional attenuation at approximately 3 [kHz], effectively blocking flexural waves in both principal lattice directions. This work demonstrates how a compact dataset and physics-informed surrogate model enable rapid, accurate inverse design of phononic metamaterials for structural noise control, avoiding laborious manual tuning or extensive parametric sweeps. By bridging plate-wave physics and deep learning, it sets a precedent for efficient AI-driven design of lightweight metamaterials with targeted vibration isolation functionality. The design of AMMs has relied mostly on the expertise of specialists, as many effects are based on localized solutions and interference. However, Gürbüz et al. [151] introduced a generative DL approach to the inverse design of acoustic metamaterials by leveraging conditional Generative Adversarial Networks (cGANs). Their objective was to generate novel unit cell geometries that achieve prescribed acoustic transmission loss characteristics over a wide frequency range (2–10 [kHz]), thereby bypassing the trial-and-error nature of traditional design methods. The authors developed a data generation framework that synthesizes pixel-based binary images of metamaterial unit cells, assigning rigid or fluid domains, and evaluated their corresponding transmission loss spectra using finite element simulations in COMSOL. A total of 2800 unit cell designs, including scatterers and locally resonant C-shaped and rectangular Helmholtz resonators, formed the dataset for training the GAN. The generator network, conditioned on a desired transmission loss spectrum and a latent vector, produced candidate cell geometries, while the discriminator distinguished between real and synthetic cell–spectrum pairs to guide learning. The trained cGAN produced diverse and acoustically meaningful unit cell designs that closely matched input transmission spectra. Simulated validation of generated designs revealed good agreement with target acoustic responses, particularly for scattering-type and Helmholtz-resonant cells, although some designs showed minor deviations in frequency or amplitude around resonant peaks. The performance was quantitatively evaluated using cosine similarity and relative error metrics, with many generated samples achieving >0.8 similarity to their target spectra. Notably, the generator was able to autonomously infer geometric features such as cavity shapes and sizes associated with specific acoustic responses, even when those features were not directly labeled, demonstrating its ability to learn implicit design rules from data. This work marks the first application of cGANs to AMM design

and establishes a blueprint for integrating machine learning with physics-based simulation. It also highlights practical considerations such as dataset balance, feature coherence, and post-processing to filter non-manufacturable geometries. The authors suggest the framework can be extended to include visco-thermal effects, finer frequency resolution, and real-world manufacturing constraints, pointing toward future applications in tailored sound insulation systems for vehicles, buildings, and industrial noise control. Mahesh et al. [152] used a deep convolutional autoencoder network for the inverse design of low-frequency acoustic absorbers, specifically targeting Helmholtz-resonator-based and hybrid absorber geometries augmented with curvilinear necks and micro-perforated panels. Leveraging an analytical model grounded in effective-medium theory and electro-acoustic analogies, the authors generated a large training dataset that maps absorber geometries to absorption spectra. The autoencoder's encoder learned a compressed representation of desired spectral responses and predicted corresponding geometric parameters, while the decoder reconstructs the absorption profile from these predicted parameters. The method was first validated on standard Helmholtz resonators and then extended to more complex hybrid absorbers, achieving high prediction accuracy without requiring prior geometric initialisation. Notably, the network architecture contained few learnable parameters, improving computational efficiency. The model is used to synthesise four distinct absorber designs that exhibited near-perfect absorption across targeted low-frequency bands (200–650 [Hz]), with designs featuring deep sub-wavelength thicknesses ($\approx \lambda/23$ at 200 [Hz]). The use of acoustic metasurfaces has emerged as a highly promising method for controlling acoustic waves due to their ultra-thin profile. However, the conventional design process for acoustic metasurfaces relies on numerical and trial-and-error methods to determine the effective properties of local resonant unit cells. This approach can be inefficient and computationally intensive, particularly when dealing with narrow cavities that require the consideration of thermo-viscous loss. To address this challenge, Donda et al. [9] introduced a DL-based inverse design framework for creating ultrathin acoustic metasurface absorbers that achieve near-perfect absorption at extremely low frequencies with thicknesses far below the typical quarter-wavelength limit. They trained a CNN on a comprehensive dataset of pixelated unit cell designs, including narrow, high-aspect-ratio cavity structures—paired with their full-spectrum absorption responses computed via thermo-viscous FEM models. This CNN effectively models the wideband absorption behavior in mere milliseconds, compared to prohibitively slow traditional simulations. Once trained, the network was used to rapidly

generate novel metasurface absorbers; the authors demonstrated an absorber with 100% absorption at just 38.6 [Hz] and thickness of only $\lambda/684$. Notably, this design greatly surpasses conventional quarter-wavelength limits, achieving a total internal path of just $\lambda/5.7$. They also combined unit cells into broadband supercells, validating both numerically and experimentally that the network-designed structures accurately produced target absorption profiles at ultralow frequencies. This work highlights the efficacy of DL in navigating complex, non-intuitive metasurface design spaces, offering a fast, highly automated, and experimentally validated route to ultra-thin, low-frequency acoustic absorbers.

Wu et al. [153] proposed a novel DL-based multi-fidelity optimisation framework to enable efficient and accurate inverse design of acoustic metasurfaces that achieve uniform scattered acoustic fields, a key requirement for applications such as acoustic cloaking. The framework integrates low and high-fidelity simulation data using a custom Multi-Fidelity Composite Convolutional Neural Network (MF-CCNN), which learns both linear and nonlinear correlations between fidelity levels. The MF-CCNN predicts the high-dimensional scattered acoustic field from physical input parameters (density and Young's modulus of metasurface units), achieving high accuracy at a significantly reduced data cost. A fully connected neural network is then used to optimize the physical parameters of the metasurface, leveraging back-propagation through the frozen MF-CCNN to minimise field variance as an objective. The method is validated through comparison with two other multi-fidelity models (MF-PCNN and BFLL-1) and single-fidelity networks. MF-CCNN consistently outperforms alternatives in terms of prediction error and generalization, with mean absolute error reduced by over 20% and optimisation time reduced to only 8% of that required by a genetic algorithm. The framework enables rapid generation of near-optimal metasurface configurations, demonstrating its strong potential for accelerating acoustic metamaterial design through deep learning and surrogate modelling.

Donda et al. [154] proposed a DL framework for both forward and inverse design of acoustic metasurface absorbers with high structural complexity and many degrees of freedom, particularly under oblique incidence conditions and in the low-frequency regime. The authors first develop a 2D CNN to predict the absorption spectra of metasurface structures based on input design parameters including a freeform 2D propagation channel, cylindrical aperture diameter, thickness, and incidence angle. To reduce computational demands and improve robustness, data augmentation and principal component analysis (PCA) are applied to the

training dataset of over 9000 COMSOL-simulated samples. The CNN model achieves accurate forward predictions with a mean squared error (MSE) below 0.01, even for ultrathin absorbers with thicknesses as low as $\lambda/64$ at 82 [Hz]. For inverse design, a cGAN is employed to generate structural patterns that yield target absorption spectra under specified conditions. The cGAN learns to produce viable metasurface designs directly from spectral requirements, offering a rapid, non-iterative alternative to traditional trial-and-error or simulation-based optimisation. The proposed approach demonstrates strong agreement between predicted and simulated absorption responses and confirms the potential of DL to accelerate and automate the design of complex, performance-driven acoustic devices.

The application of DL to the design and analysis of AMMs is still an evolving field, with growing momentum driven by advances in neural architectures, simulation fidelity, and fabrication techniques. As this convergence between computational acoustics and artificial intelligence (AI) matures, several emerging trends and future directions are reshaping the research landscape such as in the introduction of generative models, such of VAE and GAN; the use of transfer learning approaches that reuse knowledge from related tasks, incorporation of multi-scale interactions (microstructure $\rightleftharpoons$ macro-structure) and multi-physics coupling (e.g., vibro-acoustics, thermo-acoustics, or fluid-structure interactions); incorporation of experimental information into the design pipeline, use of data-augmentation to enhance DL's prediction fidelity; the integration of interpretable mechanisms into the design pipeline; and the integration with design automation and human-AI collaboration

2.7 Conclusions

The control of unwanted sound and noise has historically relied on a variety of traditional acoustic materials and configurations, each grounded in well-established physical principles. Enclosure systems, acoustic barriers, and damping treatments have long formed the backbone of noise control strategies in both industrial and architectural applications. These materials primarily function by absorbing or reflecting sound energy, and their effectiveness is heavily reliant on bulk, geometry, and placement relative to the source and receiver. Among traditional absorbers, porous materials such as mineral wool, foam, and fibrous composites have been extensively used due to their broadband absorption characteristics at mid to high frequencies. These materials operate via viscous and thermal losses as air particles oscillate within their interconnected voids. While effective in many applications, their performance

diminishes significantly at low frequencies unless substantial thickness is employed, posing limitations where space and weight are constrained. Panel sound absorbers and resonant absorbers address some of these limitations. Panel absorbers exploit membrane or plate vibrations to absorb energy at specific frequency ranges, especially in the low-frequency domain. Similarly, Helmholtz resonators, perforated panel absorbers, and micro-perforated absorbers (MPPs) use resonance effects to target narrowband absorption, often tunable through geometric parameters such as cavity volume, neck dimensions, and perforation characteristics. These systems offer more compact solutions but still suffer from trade-offs between bandwidth, tuning flexibility, and overall absorption efficiency. Despite their utility, traditional materials and approaches face inherent challenges: limited low-frequency performance without large volumes, narrowband behavior in resonant devices, and reduced adaptability to dynamic acoustic environments. These constraints have spurred the exploration of novel classes of engineered materials that enable more precise and compact control over sound propagation. Acoustic metamaterials (AMMs) represent a transformative development in this regard. Based on concepts from electromagnetic metamaterials (EMMs), AMMs are artificially structured composites designed to manipulate sound waves in ways that are not possible with conventional materials. By tailoring sub-wavelength-scale structural features, AMMs can exhibit exotic effective properties, such as negative bulk modulus or negative density, leading to phenomena such as acoustic bandgaps, negative refraction, and super-lensing. These properties enable highly selective frequency control and sub-wavelength-scale performance that traditional absorbers cannot achieve. More recently, acoustic metasurfaces, a two-dimensional counterpart of bulk metamaterials, have emerged as an even more compact and versatile platform for wave manipulation. Through judicious design of unit cell geometries and surface impedance profiles, metasurfaces allow for engineered reflection, refraction, phase modulation, and absorption, all within subwavelength thicknesses. This opens promising avenues for low-profile acoustic cloaks, flat lenses, and adaptive noise control solutions. While traditional acoustic materials have played a foundational role in noise control engineering, their limitations in efficiency, compactness, and tunability, especially at low frequencies, have driven the need for advanced engineered materials. AMMs and metasurfaces provide a fundamentally new paradigm, offering unprecedented control over sound with greater spatial efficiency and functional customisation. Furthermore, this chapter also examines the application of DL to the design and analysis of AMMs represents a paradigm

shift in how these materials are conceived, optimised, and realised. Traditional optimisation and physics-based approaches, though powerful, are often limited by computational intensity, high-dimensional parameter spaces, and design complexity. DL offers a complementary strategy, one that leverages data-driven models to capture complex acoustic phenomena, accelerate simulation workflows, and enable inverse design in ways that were previously intractable. Supervised learning methods such as CNNs, RNNs, feedforward networks, etc., have been employed for rapid performance prediction and forward modelling of AMMs. These models have demonstrated high accuracy and computational efficiency, particularly in mapping geometries to their acoustic properties. In the inverse design landscape, generative models, including VAEs, GANs, and tandem networks, have emerged as powerful tools for creating new AMM structures conditioned on specific performance targets. These methods not only reduce design time but also unlock the ability to explore vast and non-intuitive design spaces. Despite significant progress, several challenges remain, such as the need for interpretability of black-box AI models for design trust and debugging, the incorporation of robustness into the design pipeline, and the reliable integration of these methods into real-world applications. These issues will be addressed in subsequent chapters.

Chapter 3

AMM Modelling and Deep Learning Frameworks

3.1 Introduction

This chapter presents a comprehensive validation of the hybrid resonant acoustic meta-absorber (HRAM) model used in this thesis for the sole purpose of data generation for deep learning (DL) applications. In parallel, it introduces the proposed DL methodology aimed at achieving efficient and robust acoustic metamaterials (HRAM) designs. The chosen AMM structure serves as the core physical system throughout the thesis, providing the foundation for subsequent data-driven exploration and optimisation tasks. The analysis begins with the development of an analytical model of the absorber (composed of an acoustic cavity with a resonant inclusion and a top perforated plate), formulated using transfer matrix methods (TMM), which offer a fast and sufficiently accurate approximation of the system's acoustic behaviour: The TMM expresses the acoustic state variables (pressure and particle velocity) at the input and output of each constituent element through a linear matrix relationship. By cascading the matrices of individual components, the overall acoustic response of the system can be efficiently computed. Complementary finite element simulations are performed in COMSOL Multiphysics to validate the analytical predictions. To physically realise the designs, additive manufacturing using the fused deposition modelling (FDM) is employed, supporting rapid and cost-effective prototyping. The fabricated samples are experimentally characterised using a two-microphone impedance tube setup, following standardised acoustic testing protocols. Building upon the analytical model of the HRAM as the data generator, a deep learning framework termed the Structured Latent Autoencoder (SLAE) is trained (calibrate the SLAE model parameters) on this data to achieve both forward prediction and inverse inference of the absorber performance. To bridge the gap between idealised simulations and real-world measurements, data augmentation techniques are applied, incorporating known sources of manufacturing uncertainty into the training data. Finally, a transfer learning (TL) approach is employed to adapt the pre-trained SLAE model to experimental data, enabling accurate geometric inference under real-world fabrication deviations. Together, these interconnected methodological components establish a rigorous pipeline for the design, fabrication, characterisation, and modelling of AMM absorbers, laying the foundation for the results and discussions presented in subsequent chapters. Figure 3.1 shows a schematic of the overall process from design concept to robust design. This chapter is structured as follows:

Section 3.2 introduces the proposed AMM absorber configuration, outlines the overall modelling and validation workflow, and provides detailed derivations of the analytical model

based on TMF. Section 3.3 presents the numerical simulation framework, describing the FEM implementation in COMSOL Multiphysics. Section 3.4 outlines the additive manufacturing (AM) process employed to fabricate the absorber samples, including a discussion of the FDM technology used and associated manufacturing uncertainties. Section 3.5 details the experimental setup and measurement procedures used for acoustic characterisation, including the two-microphone impedance tube method adopted for all fabricated samples. Section 3.6 introduces the deep learning framework developed in this work. It covers the data acquisition and pre-processing pipeline, the deep learning model architecture, the data augmentation techniques, and the transfer learning procedure used to adapt the model to real-world experimental data. Finally, Section 3.7 summarises the key developments and methodologies presented in the chapter.




Figure 3.1: Schematic overview of the methodology adopted in this chapter, illustrating the complete workflow from AMM concept design through modelling, data-driven learning, and robust design integration.

3.2 Absorber description

The design concept chosen, based on inspiration drawn from different designs in the literature (e.g., [90, 143, 155], in this research is based on a spatially coiled Helmholtz resonator. The resulting form was designed with simplicity in mind, while maintaining sufficient complexity to investigate the accuracy of traditional modelling methods. The system consists of an

Archimedean spiral (similar to [142, 143]) embedded within an acoustic cavity, along with a top perforated plate that allows the incident wave to propagate through the device, as shown in Fig. 3.2. It is noteworthy that the coiled inclusion can take various forms, such as a rectangular or C-shaped coils, and concentric split-rings with laterally inverted openings. The spiral structure shifts the resonance to lower frequencies by elongating the effective propagation path (phase delay) within a fixed thickness, H , enabling the wave to travel a longer distance than the physical dimensions of the device. This hybrid design introduces a complex tuning framework due to the interaction of two resonance mechanisms: the acoustic cavity and the embedded spiral. The hexagonal cavity is so chosen to allow for compact and economical arrangement of units on the metasurface.

The final absorber design, based on hybrid resonant mechanism, underwent several modifications to improve analytical accuracy and facilitate additive manufacturing. Initially, a circular aperture (acoustic inlet) was used, but the analytical model was insensitive to changes in the aperture radius. Additionally, early 3D-printed samples experienced acoustic leakages when the cavity and the top plate were printed separately and then glued together. To resolve these issues, the aperture was re-designed to match the geometry of the first turn of the embedded spiral, allowing the entire absorber to be printed as a single unit, thereby preventing leakages. This adjustment also improved the analytical model by integrating the perforation into the spiral channel, ensuring geometric continuity and simplifying modelling for subsequent optimisation studies.




Figure 3.2: The unit structure of the absorber: (a) disassembled system comprising a rigid backing, hexagonal boundary and a top perforated plate, (b) parameterised spiral, and (c) assembled system

The design parameters of the absorber are presented in Table 3.1. The initial radius, r_o , although influences absorption performance, it is not directly controlled here to preserve

geometric continuity of the spiral channel. Thus, it is defined in terms of other parameters as will be described in Subsection 3.2.1

Table 3.1: Geometric design parameters defining the unit cell of the metasurface acoustic absorber. These parameters control/influence the absorption behaviour of the absorber.

Parameter	Description	Unit
L_s	Side length of the hexagonal cavity	mm
H	Depth/thickness of the cavity	mm
t_p	Thickness of the top perforated plate	mm
p	Spacing between spiral turns	mm
t	Thickness of spiral wall	mm
r_o	Initial radius of the spiral	mm
N_t	Number of spiral turns	-

3.2.1 Analytical model

Numerical models, while highly accurate and capable of capturing complex system behaviours, are often computationally intensive. This limitation becomes particularly significant in scenarios requiring extensive parametric sweeps or optimisation studies, where numerous evaluations of the system response are necessary. Although simplified numerical models (SNMs), where feasible, can significantly reduce computational cost, numerical approaches often remain inefficient in iterative design and sensitivity analyses. This is particularly true for AMMs operating at low frequencies, where narrowband, high-Q resonant behaviour necessitates fine frequency sampling to accurately resolve the system response. In addition, the presence of subwavelength features and thermo-viscous boundary layers requires local mesh refinement, further increasing computational expense.

In contrast, analytical models offer a streamlined alternative, enabling rapid evaluation of system behavior under varying conditions. These models are derived from first principles and simplified approximations, which allow for closed-form solutions or explicit expressions that can be quickly computed. As a result, they are particularly well-suited for scenarios involving high-dimensional parameter spaces or rapid evaluation of absorber designs. While analytical derivations may not fully capture all the intricacies of the system, their ability to generate rapid and sufficiently accurate predictions ensure their role as a cornerstone in the workflow of parametric and optimisation studies.

The perforation indicated by the acoustic inlet in Fig. 3.2(c) is designed to have the same cross-section as the coiled channel for analytical convenience and to this end, the initial radius, $r_o > 0$, is determined such that

$$\int_0^{2\pi} 0.5R^*(\theta)^2 d\theta - S_1 = 0, \quad (3.1)$$

where $R^*(\theta)$, defined in Eq. (3.2), is the polar equation of the spiral, and $S_1 = pH$, is the cross-sectional area of the coiled channel.

$$R^*(\theta) = r_o + \frac{(p+t)}{2\pi}\theta \quad (3.2)$$

The system is essentially a Helmholtz resonator with a spatially coiled cavity that introduces a secondary resonance channel. The channel is modeled as an equivalent straight rectangular duct with a cross-section much smaller than the working wavelength, allowing the use of narrow region acoustics to capture the thermal and viscous losses occurring within the duct. For a rectangular duct of dimensions p and H , the viscous and thermal field functions, Ψ_i , $\forall i \in \{v, h\}$, based on the low-reduced frequency (LRF) model, are given by [156, 157]

$$\begin{aligned} \Psi_i = \kappa_i^2 \sum_{m=0}^{\infty} & \left[(\alpha_m m')^{-2} \left(1 - \frac{\tan(\alpha_m p/2)}{\alpha_m p/2} \right) \right. \\ & \left. + (\beta_m m')^{-2} \left(1 - \frac{\tan(\beta_m H/2)}{\beta_m H/2} \right) \right], \end{aligned} \quad (3.3)$$

where

$$\alpha_m = \sqrt{\left(k_i^2 - \frac{2m'}{p}\right)}, \quad \beta_m = \sqrt{\left(k_i^2 - \frac{2m'}{H}\right)}. \quad (3.4)$$

$k_v^2 = -i\omega \frac{\rho}{\eta}$ and $k_h^2 = -i\omega \frac{\rho C_p}{\kappa}$, are the viscous and thermal wave numbers, respectively, and $m' = (m + 1/2)\pi$. The upper limit of the infinite series in Eq. (3.3) is chosen such that the acoustic response is stable. C_p , ρ , η and κ are the specific heat capacity at constant pressure, the density of air, the dynamic viscosity of air and the thermal conductivity of the fluid, respectively. The complex wave number, k_c , and the complex density of air, ρ_c , in the medium can be derived as

$$\begin{aligned} k_c^2 &= k^2 \left(\frac{\gamma - (\gamma - 1)\Psi_h}{\Psi_v} \right), \\ \rho_c &= \frac{\rho}{\Psi_v}, \end{aligned} \quad (3.5)$$

where k is the wave number. The transfer matrix of each region can be expressed as

$$T_j = \begin{bmatrix} \cos(k_j L_j) & iZ_j \sin(k_j L_j) \\ i \frac{\sin(k_j L_j)}{Z_j} & \cos(k_j L_j) \end{bmatrix}, \quad (3.6)$$

where L_j , k_j and Z_j , $\forall j \in \{\text{channel, cavity}\}$, are the length, wave number and normalized impedance, respectively, associated with each region, $Z_{\text{channel}} = \frac{\rho_c c_c}{S_1}$ and $Z_{\text{cavity}} = \frac{\rho_c}{S_2}$ (c_c is the complex speed of sound in the channel). It is noteworthy that the coiled channel occupies part of the cavity space so that the effective volume of the cavity is reduced. The effective depth of the cavity is $L_{\text{cavity}} = v_f H$, where V_{sp} is the volume of the spiral and is given by

$$V_{sp} = H \int_{2\pi(N_t-1)}^{2\pi N_t} 0.5(R(\theta) + t)^2 d\theta. \quad (3.7)$$

The channel length is given by $L_{\text{channel}} = L_1 + L_2$. The length, L_2 , is computed along the centroid of the channel, as

$$L_2 = \int_0^{2\pi(N_t-1)} \left(\sqrt{R^*(\theta)^2 + \left(\frac{d}{d\theta} R^*(\theta) \right)^2} \right) d\theta, \quad (3.8)$$

noting that r_o is effectively modified as $r_o + t + 0.5p$ in this case, and $L_1 = H + tp$. An alternative approach to computing the effective length through the aperture is to account for its coiled geometry. Thus, L_1 may be re-written as: $L_1 = \sqrt{(2\pi r_a)^2 + tp^2} - \zeta$, where $r_a = r_o + (p + t)/2$, is the effective radius of the aperture, obtained by averaging $R^*(\theta)$ over the first turn, and $\zeta = 0.96\sqrt{S_1}$ [158]. This alternative is found to provide a more accurate result over a wider operational range. The overall transfer matrix of the single unit can thus be obtained as $T = T_{\text{channel}} \cdot T_{\text{cavity}}$, and the scaled surface impedance is defined as $Z = S_2 \cdot T_{1,1}/T_{2,1}$, where $T_{1,1}$ and $T_{2,1}$ refer to the elements in the first column of matrix T , and $S_2 = 1.5\sqrt{3}Ls^2$ is the cross-sectional area of the incident unit surface. The complex reflection coefficient of the unit, R , and the absorptivity, α , are subsequently calculated according to Eq. (3.9).

$$R = \frac{Z - Z_o}{Z + Z_o}, \quad \text{and} \quad \alpha = 1 - |R|^2, \quad (3.9)$$

where $Z_o = \rho c$ is the impedance in air. For a perfect absorption ($\alpha \approx 1$), the impedance matching condition requires that the surface impedance of the absorber matches that of the background medium, i.e., $\text{Re}(Z) = Z_o$, and $\text{Im}(Z) = 0$.

3.2.2 Broadband model

Broadband absorption has garnered significant attention in the field of noise control due to its practical applications. However, one of the primary challenges in designing broadband absorbers lies in achieving an optimal balance between compactness and wideband performance. While thicker absorbers generally excel at low-frequency sound attenuation, space constraints often necessitate innovative approaches, such as locally resonant structures [159, 160] or metamaterial-based solutions [9, 161–163]. Recent advancements in computational modeling and optimisation, particularly the integration of machine learning and deep learning techniques, have further accelerated the design process. These approaches enable the rapid prototyping and testing of novel absorber geometries, paving the way for more efficient and compact solutions. The narrowband design is extended here for the broadband case. To broaden the typically narrow effective absorption bandwidth of the absorber, four (4) inhomogenous units are parrallely combined to form a single unit as schematised in Fig. 3.3. Each unit of the supercell contributes six (5) parameters, so that the supercell comprises twenty (24) parameters in total. The parameter, Ls , is same for all the component cells to facilitate arrangement, while H may vary for each cell. t and tp are again held constant as with the narrowband case. The total impedance of the supercell is obtained using the electro-acoustic theory as in Eq. (3.10).

$$\frac{S_t}{Z_t} = \sum_{i=1}^4 \frac{S_i}{Z_i}, \quad (3.10)$$

where $S_i = S_2$ ($\forall i \in \{1, \dots, 4\}$), and $S_t = \sum_{i=1}^4 S_i$, is the surface area of the supercell. Z_i is the impedance of the i -th unit, and Z_t is the total impedance of the supercell.




Figure 3.3: (a) Schematic of the supercell composed of four inhomogeneous units arranged parallelly, in an edge-to-edge hexagonal tiling. While four units are considered in this thesis, the configuration can be generalised to include more units as required by the application.

3.3 Numerical model setup

Numerical modeling based on finite element analysis (FEA) is performed in COMSOL[®] Multiphysics software. The thermo-viscous acoustic module is utilised to account for the thermal and viscous losses occurring within the narrow channels of the absorber. An acoustic manifold with thickness, $2(H+tp)$, is defined at the top of the absorber to simulate the surrounding environment and an incident sound wave, defined by a background pressure field with unit amplitude, normally impinges on the absorber surface in the $-z$ direction. Furthermore, a perfectly matched layer is defined above the hexagonal volume to simulate an open and non-reflecting infinite domain. Hard boundaries are imposed on the interfaces between air and solid due to the high impedance mismatch between air and solid materials [164]. The absorber is discretised using a tetrahedral element with a maximum element size of $\lambda/20$, where λ is the wavelength corresponding to the highest frequency of interest, while the minimum element size is set to $\delta/10$, where δ is the thickness of the boundary layer defined as

$$\delta = 220[\mu m] \sqrt{100[Hz]/f_o}. \quad (3.11)$$

The value $220 [\mu m]$ in Eq. (3.11), is the thickness of the viscous boundary layer in air at $100 [Hz]$, $20^\circ C$ and $1 [atm]$, and f_o is the maximum study frequency. The perforation is meshed with a fine resolution. Boundary layers with a first layer thickness of 0.4δ are defined at the air-solid interfaces of the absorber. These settings, determined through experimentation, are applied consistently across all subsequent simulation studies. A separate mesh study

in this work further refines the mesh resolution in the perforation region, as previous settings proved less accurate at extremely low frequencies, albeit with a higher computational cost. Notably, this refinement is applied only to the perforation, where the bulk of acoustic dissipation occurs. All numerical simulations are performed on a Precision 5820 Tower workstation equipped with a 12-core Intel® processor running at 3.5 GHz and 256 GB of RAM.

3.3.1 Linearised Navier-Stokes equations

The losses within the narrow coiled path of the absorber system has been captured by the thermo-viscous acoustic module of COMSOL, which is based on solving the the Linearised Navier-Stokes equations expressed in time harmonic form $\exp j\omega$. These equations are formulated in the frequency domain to capture losses owing to viscosity and thermal conduction in an acoustic system. Acoustic variations in pressure, velocity, and temperature are thus computed in the presence of any stationary isothermal or non-isothermal background mean flow.

The linearised Navier-Stokes equations in the frequency domain interface are the linearised continuity, momentum (Navier-Stokes), and energy equations. The physics interface solves for the acoustic variations in the pressure, p , velocity field u , and temperature, T , using a first-order perturbation around the steady-state background condition, i.e.,

$$p = p' + p_o; \quad u = u' + u_o; \quad T = T' + T_o \quad (3.12)$$

where p_o , u_o and T_o are background conditions, p , u , and T are the total variables and p' , u' and T' , are the acoustic (perturbed) variables. The constitutive equations for the stress tensor and the linearised equation of state (density perturbation) in the frequency domain are given by:

$$\sigma = -pI + \mu(\nabla u + (\nabla u)^T) + (\mu_b - \frac{2}{3}\mu)(\nabla \cdot u)I \quad (3.13)$$

$$\rho = \rho_o(\beta_T C_p - \alpha_p T) \quad (3.14)$$

where β_T is the isothermal compressibility, α_p is the coefficient of thermal expansion, and C_p is the specific heat capacity (heat capacity per unit mass) at constant pressure. At the fluid-solid boundary, a no-slip isothermal condition is applied i.e., $u = 0$, and $T = 0$. The fluid simulation parameters used in this study are presented in Table 3.2

Table 3.2: Fluid simulation parameters.

Symbol	Value	Unit
T_o	343.2036	K
p_o	101325	Pa
ρ_o	1.2	kg/m^3
μ	1.825e-5	Pa.s
μ_b	0	Pa.s
C_p	1005.4	J/kgK
c	343	m/s

3.4 Additive manufacturing

Recent advances in additive manufacturing (AM), commonly referred to as 3D printing, have significantly transformed the landscape of rapid prototyping and functional material design. What was once limited to digital simulations and CAD visualisations can now be physically realised with high precision and at low cost. This evolution has been particularly impactful in the field of AMMs, where the ability to produce intricate, sub-wavelength geometries is critical for achieving desired wave manipulation effects. Complex geometries that were only digitally realisable can now be manufactured rapidly with precision. One of the most compelling advantages of 3D printing is the capacity for rapid iteration. Designs can be quickly modified, fabricated, and tested without the lead time and expense traditionally associated with subtractive or injection-molding processes. This flexibility facilitates not only research and development but also the customisation of parts for niche applications. Several 3D printing technologies exist each with its own challenges and their impacts on additively manufactured parts. Factors such as surface roughness, and dimensional fidelity, can influence how accurately a printed metamaterial performs compared to its simulated model. These effects must be considered carefully during the modeling and validation process [17, 21].

In this thesis, fused deposition modelling (FDM) is used as the primary fabrication technique. While FDM is widely used in prototyping due to its cost-effectiveness, it also introduces challenges that must be addressed. Layer adhesion variability, surface roughness, and dimensional fidelity, are among the factors that can influence the performance of fabricated acoustic structures. To ensure that the influence of FDM-induced variability is not neglected, a dedicated focus is placed on quantifying and compensating for these effects in later chapters. Specifically, DL techniques are employed to integrate manufacturing-related uncertainties into the forward and inverse modeling of AMMs. By incorporating these imper-

fections through experimental samples into the learning models, the system can be trained to understand how variations in print quality and dimensional accuracy affect the acoustic response, enabling more robust and realistic predictions

All the samples used in this work are fabricated using polylactic acid (PLA) on a Prusa i3 MK3S+ 3D printer. The CAD models are exported in .stl format and processed using PrusaSlicer. Default slicing settings are used, with the exception of the temperature: the first layer was printed at 215 °C to enable it stick very well to the build plate and prevent warping or detachment during printing, while all subsequent layers are printed at 200 °C to reduce stringing or oozing.

3.5 Experimental setup

The test rig used in this research work is a cylindrical impedance tube of diameter, $d = 127$ [mm]. A high-amplitude loudspeaker, housed in a custom-built plywood unit with a circular flange opening for mounting onto one end of the aluminium tube, is attached to the tube at one end. The loudspeaker is a JBL 2206H/J12 [in.] sub-woofer. An adjustable back-cavity depth termination/sample holder is built around the other end of the tube to accommodate different sample thicknesses, and allows for samples to be printed on a 199.5×199.5 [mm] square panel. It has a depth of 65 [mm] and a 24.75 [mm] wall thickness. Two G.R.A.S. 40 PH microphones, located at positions $x_1 = 0.212$ [m] and $x_2 = 0.127$ [m] from the end of the tube, are used to instrument the rig in accordance with ISO 10534-2:2001 [165]. The tube schematic setup is shown in Fig. 3.4. Two microphones, M_1 and M_2 , are located at positions x_1 and x_2 , respectively, and white noise propagates from the loudspeaker through the tube. Using the measured transfer function between M_1 and M_2 , H_{12} , to correct for mismatch in their frequency responses via the switching technique [165], the complex reflection coefficient under plane wave assumption is given by Eq. (3.15)

$$R = \frac{H_{12} - e^{-jks}}{e^{jks} - H_{12}} e^{2jkx_1}, \quad (3.15)$$

where $s (= x_1 - x_2)$ is the microphone spacing. The absorptivity is subsequently calculated from $\alpha = 1 - |R|^2$. The lower and upper frequency limits of the measurement are functions of the microphone spacing and tube diameter, respectively. Thus, the lower limit of the tube is $0.01c/s$ [166] (≈ 40 [Hz]), and the upper limit is $0.586c/d$ (≈ 1568 [Hz]). The frequency resolution used in the experiment is approximately 0.6 [Hz] with a sample rate of 40,000 [Hz].




Figure 3.4: Schematic of the two-microphone impedance tube setup used for acoustic absorption measurements. Adapted from Hossain et al. [10].

Figure 3.5 shows the laboratory setup of the impedance tube during a testing session of the acoustic absorber.




Figure 3.5: Photograph of the experimental impedance tube setup used for acoustic absorption measurements in a typical testing session. Key components include the loudspeaker, two calibrated microphones, sample holder, and data acquisition system.

3.6 Deep learning framework

One benefit of the DL framework is to accelerate the on-demand design process for AMMs. Traditional optimisation methods can be computationally intensive, especially when dealing with complex designs with high dimensionality [167]. DL offers a more efficient alternative by enabling near-instantaneous predictions (after training), often in the order of milliseconds [167], which makes real-time design iteration and inverse design feasible. Moreover, this framework is designed to integrate manufacturing imperfections, such as fabrication tolerances, surface roughness, or material inconsistencies, into the design process itself. Rather than treating these imperfections as post-fabrication issues, the proposed approach embeds them into the training data and learning objectives, allowing the model to predict and optimize for performance under more realistic, imperfect conditions. Thus, the methodology is outlined as follows.

3.6.1 Data acquisition

Often, for any deep neural network based technique, a large dataset is essential to obtain a good prediction model. Here, the generation of sufficient data using experimental or numerical means is practically a challenging task, and thus the training data for the DL model are generated using the analytical model of the absorber described in Subsection 3.2.1. The input to the network comprises the spectral characteristics, i.e., $[\alpha_p, f_r, f_u, f_L]$, while the outputs are the corresponding geometric parameters, p and Nt . It is noteworthy that H and L_s are coalesced in the input spectra to provide the designer with the flexibility to tailor the device's size as needed on the one hand, and to address the non-uniqueness problem of inverse mapping on the other hand. To this end, the input to the network becomes $[\alpha_p, f_r, f_u, f_L, L_s, H]$, where $t=0.8$ [mm], and $t_p = 1.5$ [mm] are fixed during modelling. The band-edge frequencies, f_u, f_L , are derived from the spectral data using the full width at half maximum (FWHM) method. This approach identifies the range of frequencies over which the absorptivity remains above half of its peak value.

Let $\alpha(f)$ denote the absorption coefficient as a function of frequency, and let α_p be the peak absorption. The half-maximum threshold is defined as:

$$\alpha_{half} = \frac{\alpha_p}{2}. \quad (3.16)$$

The frequencies at which $\alpha(f) = \alpha_{half}$ are then computed on both sides of the peak, the lower bandwidth frequency, f_L , and the upper bandwidth frequency, f_u . The bandwidth, Δf , is defined as the difference between these two frequencies, i.e., $\Delta f = f_u - f_L$. This method captures the effective spectral width where the absorber maintains strong performance, and is widely used in acoustic and electromagnetic metamaterials to characterise resonance-based behavior [168–171].

In many DL approaches for AMM design found in the literature, the networks were typically trained on the full absorption spectrum. While this provides comprehensive information, it also leads to high input dimensionality, more complex network architectures, and increased computational costs, which make training and hyperparameter tuning more difficult. In contrast, this thesis adopted a more compact and efficient strategy by training on the most critical features of the spectrum. This significantly reduces the dimensionality of the input space, enabling faster training, simpler network architectures, and reduced computational complexity, while still maintaining sufficient accuracy for design prediction.

The geometric parameters were randomly sampled from a uniform distribution over their respective ranges, as presented in Table 3.3, ensuring adequate coverage of the design space while respecting experimental constraints. A total of 250,000 samples are generated with a resolution of 0.5 [Hz], and divided into 90% training and 10% hold-out sets. The training set was further divided into 70% training and 30% validation sets as illustrated in Fig. 3.6.




Figure 3.6: Data distribution scheme used for deep learning training and validation. The diagram illustrates how the dataset is partitioned into training, validation, and test subsets, supporting model development and performance evaluation

Table 3.3: Parameter ranges for deep learning, where $t = 0.8$ [mm] and $tp = 1.5$ [mm] are fixed

Parameter	Range	Unit
p	1 - 15	mm
Nt	1.5 - 5	-
Ls	10 - 50	mm
H	5 - 30	mm

The hold-out test set is used to independently evaluate the performance of the network after the training process is completed, ensuring an unbiased assessment of generalisation capability. Meanwhile, the validation set is employed during training to monitor the network's performance and to determine when to trigger the stopping criteria. This criteria is triggered when the performance on the validation set deteriorates or does not improve for a consecutive number of training epochs, to prevent overfitting. In this study, an early stopping threshold of 50 epochs was used.

To enhance the learning process and ensure numerical stability during training, both the input and output data are normalized according to Eq. (3.17).

$$u_n = \frac{u - u_{\min}}{u_{\max} - u_{\min}}, \quad (3.17)$$

where u_n represents the normalised value, u is the original, unnormalised data point, $u_{\min}$ and $u_{\max}$, are the minimum and maximum values, respectively, of the feature (variable). The transformation ensures that all features contribute equally to the learning process and prevent dominance by variables with larger magnitudes. It is noteworthy that, although the absorption spectra already have normalised values, the size-related features (H and Ls) included in the spectra varies significantly over wider ranges, thus normalising the features together ensures a uniform transformation process across the dataset. This avoids introducing inconsistencies where some variables are processed differently than others. In order to account for real physical constraints in the acquired data, the geometric constraints in Eq. (3.18) are imposed during data generation to ensure non-negative values of r_o and to prevent the embedded coil from intersecting or spiraling into the walls of the cavity in any sample instance.

$$r_o \geq 0; \quad \frac{Ls}{R(\theta)|_{\theta=2\pi N_t} + t} > 1, \quad (3.18)$$

where

$$R(\theta) = r_o + \frac{(p+t)}{2\pi}\theta. \quad (3.19)$$

3.6.2 Model architecture

Considering the significantly reduced dimensionality of the input space, a network with moderate complexity is sufficient. Thus, a structured latent autoencoder (SLAE) is implemented, where the latent space is explicitly structured to learn the geometric features of the absorber. An autoencoder network consists of two main components: an encoder, which compresses the input into a lower-dimensional latent representation, and a decoder, which reconstructs the original input from this representation. Let the input vector be denoted as:

$$\mathbf{u} = [\alpha_p, f_r, f_u, f_L, Ls, H] \in \mathbb{R}^6 \quad (3.20)$$

where the first four (4) components represent key spectral features of the acoustic absorber, and the last two are size-related parameters for customization purposes. The encoder function, $\mathbf{E}$, maps the input vector, $\mathbf{u}$, to a latent representation $\mathbf{z} \in \mathbb{R}^m$, with m being the dimension of the latent space:

$$\mathbf{z} = \mathbf{E}(\mathbf{u}) \quad (3.21)$$

Here, the latent space is supervised to learn the absorber geometric properties, namely p and N_t . The parameters Ls and H , are directly injected into the latent representation to enhance reconstruction fidelity. The decoder function, $\mathbf{D}$, then reconstructs the spectral characteristics:

$$\hat{\mathbf{u}} = \mathbf{D}(\mathbf{z}) \quad (3.22)$$

Loss function

The model is trained using a combination of supervised and reconstruction losses. The primary loss components are:

- **Reconstruction loss:** Measures how well the decoded output, $\hat{\mathbf{u}}$, matches the original spectral input $\mathbf{u}$. Mean squared error (MSE) is used:

$$\mathcal{L}_{\text{recon}} = \frac{1}{N} \sum_{i=1}^N \left\| \hat{\mathbf{u}}^{(i)} - \mathbf{u}^{(i)} \right\|^2 \quad (3.23)$$

- **Latent supervision loss:** Ensures that the latent vector, $\mathbf{z}$, match the true geometric parameters $\mathbf{g} \in \mathbb{R}^2$. This encourages the network to embed physical meaning into part

of the latent space:

$$\mathcal{L}_{\text{geom}} = \frac{1}{N} \sum_{i=1}^N \left\| \mathbf{z}_g^{(i)} - \mathbf{g}^{(i)} \right\|^2 \quad (3.24)$$

The total loss is a weighted sum of both components:

$$\mathcal{L}_{\text{total}} = \psi \cdot \mathcal{L}_{\text{recon}} + \beta \cdot \mathcal{L}_{\text{geom}}, \quad (3.25)$$

where ψ and β are tunable hyper-parameters that control the balance between reconstruction accuracy and supervision of the latent geometry, and are set to 0.005 and 1, respectively. The SLAE enables a partial embedding of physical meaning into the latent space by explicitly supervising the first two latent dimensions to represent geometric design parameters. The network architecture, illustrated in Fig. 3.7, consists of two main components: the encoder and the decoder. The encoder maps the 6D input vector to a 4D latent representation, through a sequence of four fully connected layers with dimensions $6 \rightarrow 64 \rightarrow 128 \rightarrow 64 \rightarrow 4$. ReLU activation functions are applied after each hidden layer, while the final latent layer uses linear activation. Layer normalization is applied after each hidden layer (excluding the latent layer) to stabilize training and accelerate convergence. The decoder adopts a symmetric architecture with layer sizes $64 \rightarrow 128 \rightarrow 64 \rightarrow 6$, to reconstruct the original input from the latent space. The full spectrum may be recovered from predicted spectral features using the Lorentzian model defined in Eq. (3.26).

$$L(f) = \alpha_p \cdot \frac{\gamma^2}{(f - f_r)^2 + \gamma^2}, \quad (3.26)$$

where $\gamma = \frac{f_u - f_L}{2}$, and f is the frequency grid vector which is defined over the range 50 - 500 [Hz] throughout this thesis. It is important to note that the decoder component of the network was used solely for validation purposes and is not needed for inverse (geometry) predictions in practice.

Network training is performed using the ADAM optimizer, with a learning rate of 1×10^{-4} , weight decay of 1×10^{-5} , and a batch size of 64 samples, to minimize the loss in Eq. (3.25). To prevent overfitting and unnecessarily long training epochs, early stopping is employed with a stopping threshold of 50 epochs. The performance of the network is evaluated based on mean squared error (MSE), and coefficient of determination (R^2) in Eq. (3.27).




Figure 3.7: Schematic of the proposed SLAE architecture, consisting of an inverse model (encoder) that predicts a 2D latent vector of geometric features from the given spectral features, and a forward model (decoder) that reconstructs the original input features from the latent space. The size-related features (L_s and H) are directly passed into the latent space to support reconstruction. The encoder comprises 4 fully connected layers, including the latent layer, with ReLU activations, while the decoder mirrors this architecture. Both modules incorporate layer normalization (LN).

$$R^2 = 1 - \frac{SS_{\text{residual}}}{SS_{\text{total}}}, \quad (3.27)$$

where:

- $SS_{\text{residual}} = \sum_{i=1}^n (u_i - \hat{u}_i)^2$, is the residual sum of squares, representing the sum of squared differences between the observed values, u_i , and the predicted values, $\hat{u}_i$.
- $SS_{\text{total}} = \sum_{i=1}^n (u_i - \bar{u})^2$, is the total sum of squares, representing the variance of the observed data, calculated as the sum of squared differences between the observed values, u_i , and their mean, $\bar{u}$, and
- n , is the number of data samples.

3.6.3 Data Augmentation

In practical AMM design workflows, training data is often obtained from numerical/analytical simulations based on idealised models. These simulations do not capture the imperfections and uncertainties introduced during real-world fabrication processes such as 3D printing. As a result, models trained solely on simulated data may fail to generalise well to experimental or fabricated prototypes. A well-trained model should make predictions that are invariant to spectral variations of the same geometry. One way to achieve this is data augmentation, where the training set is artificially augmented by including random perturbations of the input samples under different transformations that preserve the geometry (or labels) [172].

Table 3.4: Summary of the proposed SLAE architecture, showing the encoder–decoder layer types, output shapes, and number of trainable parameters

Layer type	Output Shape	No of trainable parameters
Linear	[-1, 64]	448
LayerNorm	[-1, 64]	128
ReLU	[-1, 64]	0
Linear	[-1, 128]	8,320
LayerNorm	[-1, 128]	256
ReLU	[-1, 128]	0
Linear	[-1, 64]	8,256
LayerNorm	[-1, 64]	128
ReLU	[-1, 64]	0
Linear	[-1, 4]	260
Linear	[-1, 64]	320
LayerNorm	[-1, 64]	128
ReLU	[-1, 64]	0
Linear	[-1, 128]	8,320
LayerNorm	[-1, 128]	256
ReLU	[-1, 128]	0
Linear	[-1, 64]	8,256
LayerNorm	[-1, 64]	128
ReLU	[-1, 64]	0
Linear	[-1, 6]	390

Data augmentation is extensively utilised in neural network-based pattern recognition tasks [173–176], and techniques such as translation, deformation, and reflection [173, 176], have significantly enhanced accuracy in image recognition problems. In recent years, DL has made dramatic impact in AMM modeling and has delivered state-of-the-art performances. However, work related to data augmentation for deep learning in AMM research has been rarely reported. In this thesis, random transformations are programmatically applied to the spectral responses to enhance data diversity and, more importantly, to mitigate the mismatches observed in the SLAE, which originated from analytical data. Care is taken to ensure that the applied augmentations preserve the underlying absorber physics by avoiding excessively large or non-physical alterations.

The impact of fabrication and imperfections in theoretical models on the acoustic response is investigated through the following spectral transformations: bandwidth modulation, spectrum scaling, spectrum shifting, and spectrum noise. In this phase, data augmentation is applied directly to the four key spectral features. This approach simplifies both augmentation and normalization processes, and facilitates easier interpretation during inference. Given the 4 spectral feature vector $\mathbf{u} = [\alpha_p, f_r, f_u, f_L]$, the following transformations are applied.

3.6.4 Bandwidth modulation

Bandwidth modulation adjusts the effective bandwidth of the spectral response by scaling the distances between the peak frequency and the half-power frequencies. Given the right and left frequency distances relative to the peak frequency as

$$d_u = f_u - f_r, \quad d_L = f_L - f_r. \quad (3.28)$$

A random scaling factor, s , is sampled uniformly from the range $s \sim \mathcal{U}(s_{min}, s_{max})$. The modulated half-power frequencies are then calculated

$$\begin{aligned} f'_u &= f_r + s \cdot d_u, \\ f'_L &= f_r + s \cdot d_L. \end{aligned} \quad (3.29)$$

f_r and α_p , remain unchanged.

3.6.5 Peak scaling

To simulate variations in spectral intensity, the peak absorption, α_p , is scaled by a random factor k , sampled from $k \sim \mathcal{U}(k_{min}, k_{max})$. The scaled peak is computed as

$$\alpha'_p = \alpha_p \cdot k. \quad (3.30)$$

The other feature values remain unchanged during this transformation.

3.6.6 Frequency shifting

Frequency shifting translates the frequency components uniformly to simulate effects such as temperature variation or sensor calibration errors. A random frequency offset, Δf , is drawn from $\Delta f \sim \mathcal{U}(-\delta, \delta)$, where δ controls the maximum shift magnitude (e.g., $\delta = 0.007$ corresponding to a shift of 3 [Hz]). The shifted frequencies are given by

$$\begin{aligned} f'_r &= f_r + \Delta f, \\ f'_u &= f_u + \Delta f, \\ f'_L &= f_L + \Delta f. \end{aligned} \quad (3.31)$$

3.6.7 Noise addition

To mimic environmental or measurement noise, additive Gaussian noise, $\mathbf{n}$, is applied element-wise to the feature vector, $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I})$, where σ_n is a small standard deviation, and $\mathbf{I}$ is the identity matrix. The noisy feature vector is then written as

$$\mathbf{x}' = \mathbf{x} + \mathbf{n}. \quad (3.32)$$

This feature-level augmentation approach provides computational efficiency and preserves key physical characteristics, improving model generalization.

This feature-level augmentation approach provides computational efficiency and preserves key physical characteristics, improving model generalization while avoiding complex full-spectrum transformations.

Figure 3.8 illustrates sample augmentation scenarios. Augmented samples are generated dynamically on-the-fly using the described transformations, applied sequentially within each training batch, with 50% split between real and augmented samples per batch. Noise is added with a probability of 20% and a noise factor of 1%. Informed by previous validation studies [177], deviations in absorption peak amplitude are constrained to within $\pm 10\%$ and biased toward amplitude reduction. The maximum frequency shift is limited to ± 3 [Hz]. To improve network robustness, a broader augmentation range is allowed during training (up to ± 4 [Hz]), this larger range was used only for training augmentation, it is not presented as an observed physical shift. Bandwidth augmentation is applied within the range [1.0, 3.0], allowing for bandwidth expansion only. It is noteworthy that the augmentations applied in this work are not intended to rebalance the data distribution or to synthesize new geometric labels, as generating physically consistent target parameters for augmented spectra is not feasible under the imposed geometric constraints. Instead, the proposed augmentation strategy introduces controlled, physically motivated variability to emulate known discrepancies between analytical predictions and experimentally measured responses. These discrepancies arise from cumulative effects such as manufacturing tolerances, surface roughness, and unmodeled viscothermal losses, which predominantly affect the spectral response rather than the nominal geometry. Consequently, all augmented samples retain the target parameters of their corresponding original prototypes during training. This formulation reflects the practical reality that a single manufactured geometry may exhibit a range of plausible acoustic responses, rather than a unique deterministic spectrum

3.6.8 Transfer learning

The 3D printing technology fabricates components by building them up layer by layer using engineering materials, as opposed to subtracting material, which helps minimize waste while achieving a high-quality surface finish [17, 21]. In addition, it reduces production costs by eliminating the need for traditional manufacturing tools and has seen significant growth in both its applications and production volume. Several 3D printing technologies exist [17], but the most widely used is FDM. This method operates by employing a solid plastic filament, typically PLA, as the primary material. During the process, the filament is fed into the machine, where it is heated to a temperature sufficient to melt the plastic. The melted material is then extruded through a nozzle in a controlled manner. Layer by layer, the machine deposits the molten filament onto a build platform, with each layer adhering to the one below it as it cools and solidifies. This sequential layering continues until the desired 3D object is fully formed. A significant drawback of FDM is the poor surface finish often observed in the final products. This issue arises primarily from the layer-by-layer deposition inherent to the fabrication process, which creates visible layer lines. Furthermore, the tessellation of the original CAD model significantly contributes to surface irregularities, as the geometric approximation of curves and surfaces impacts the smoothness of the printed specimen [18]. To predict or address surface roughness, researchers have explored methods such as integrating informed roughness factors into numerical models [18, 178] or utilising machine learning techniques [179–183]. While machine learning approaches offer state-of-the-art predictions, further work is required to effectively integrate the predicted information into realistic model development.

The lack of consistency in part quality and reliable repeatability of the manufactured part [17, 21] stems from the intricate and dynamic nature of the manufacturing process, where numerous variables, such as material properties, geometric features, and processing conditions, interact in complex ways. Even within a single system, capturing the underlying physics is a daunting task due to these constantly changing factors [184]. Consequently, the manufacturing process can unintentionally alter the behavior of printed parts or compromise the structural integrity of the material [17, 18, 21, 185–187]. This discrepancy often arises due to inherent process-induced defects such as geometrical deviations, porosity, surface irregularities, surface roughness, and voids, among others [17, 188]. Consequently, the

obtained material behavior may deviate significantly from the intended design, which can lead to suboptimal or undesirable performance in applications.

However, while the influence of additive manufacturing is often viewed as a drawback in this regard, it can also be leveraged as an advantage. When these process-induced imperfections are strategically controlled or harnessed, they can serve as a tool to achieve targeted material behaviours that may not otherwise be feasible through conventional manufacturing methods. For example, controlled porosity can be used to enhance acoustic absorption or thermal insulation properties. In the case of biological implants, controlling surface-breaking pores allows for better integration with biological tissues [189], while surface irregularities may be exploited to improve mechanical interlocking/adhesion or to scatter waves for uniform energy distribution. Similarly, geometrical deviations in periodic structures can be used to fine-tune the dynamic response of metamaterials, enabling unique mechanical or acoustic functionalities.

Such an approach requires a paradigm shift in how manufacturing defects are perceived, not as flaws to be eliminated in many applications, but as design opportunities to be optimised. By integrating computational modelling, advanced characterisation techniques, and DL-based process optimization, it is possible to intentionally introduce or tailor these imperfections to enhance the performance of advanced materials and devices. This defect engineering approach holds great promise, particularly for fields such as AMMs, where the interaction between structure and wave behavior is highly sensitive to small-scale variations. While the imperfections inherent in additive manufacturing processes can pose challenges, they also offer a unique opportunity to realize materials with customized and application-specific properties. By shifting focus toward controlling and utilizing these effects, one can turn manufacturing-induced defects into a “blessing in disguise”, opening new frontiers in material design and functionality.

The central objective of this thesis is to extend the predictive capability of a modelled system beyond ideal simulations to real-world fabricated absorbers. To this end, transfer learning (TL) is employed to incorporate experimental measurements and adapt the pre-trained SLAE to handle manufacturing imperfections. TL is a machine learning paradigm where knowledge learned from one domain (source domain) is re-used or adapted to improve performance on a different but related domain (target domain). In this context, the source domain comprises synthetic data generated from an analytical model, while the target do-

main consists of spectral samples obtained from experimental measurements. Fabrication techniques such as FDM used in this thesis, introduce imperfections including surface roughness, voids, dimensional inaccuracies, and material property variation, which can significantly affect the absorption performance. These factors are typically not accounted for in simulated data. Therefore, a naive model trained only on synthetic data will not be robust in practice.

To achieve this, a small set of experimental data is introduced to the pre-trained SLAE model, and the network is fine-tuned using TL. Specifically, the encoder and decoder weights are frozen—except for the last two layers—to preserve the learned spectral embedding. These final layers remain trainable and are re-trained using the experimental data. This approach enables the network to adapt to real-world artifacts while retaining the general knowledge acquired from simulation. Figure 3.9 illustrates the TL process.

Let θ_{sim} denote the weights learned from simulation, and θ_{exp} the updated weights after fine-tuning on experimental data. The objective is to minimize the following adapted loss on the experimental set, $\mathcal{D}_{\text{exp}} = \{(\mathbf{u}, \mathbf{g})\}$:

$$\min_{\theta_{\text{exp}}} \mathbb{E}_{(\mathbf{u}, \mathbf{g}) \sim \mathcal{D}_{\text{exp}}} [\mathcal{L}_{\text{total}}(\mathbf{u}, \hat{\mathbf{u}}; \theta_{\text{exp}})] \quad (3.33)$$

where $\mathcal{L}_{\text{total}}$ is the same combined loss function defined in Eq. (3.25), $\mathbf{g}$ is the ground truth geometric parameters, and $\mathbf{u}$ is the spectral input to the model. The expectation operator, $\mathbb{E}_{(\mathbf{u}, \mathbf{g}) \sim \mathcal{D}_{\text{exp}}}$, denotes the average of this loss over all samples drawn from the empirical distribution defined by the experimental dataset. By initializing $\theta_{\text{exp}} \leftarrow \theta_{\text{sim}}$, the network benefits from prior knowledge and requires fewer experimental samples to converge to an accurate solution. This TL strategy ensures that the model can anticipate and compensate for discrepancies introduced during fabrication, resulting in a more robust and practically deployable predictor of AMM performance. The TL workflow is shown in Fig. 3.10.




Figure 3.8: Illustration of sample data augmentation scenarios. The red curve represents the original (theoretical) spectrum, while the black curves depict augmented spectra generated using various transformation techniques. These transformations include, (a) bandwidth adjustment to reflect potential broadening or narrowing of spectral features, (b) peak scaling to emulate variations in absorption amplitude, (c) horizontal shifts to simulate frequency drift, and (d) noise addition to account for experimental imperfections. Each augmented spectrum demonstrates the ability of the applied methods to introduce realistic variability while preserving key spectral characteristics.




Figure 3.9: Schematic representation of the transfer learning process. Only the final two layers of the encoder and decoder, highlighted in red, are trainable, while all other layers remain frozen. The network input consists of experimentally measured spectra that reflect inherent manufacturing imperfections.




Figure 3.10: Overview of the design workflow for the robust AMM design. The pipeline incorporates a pre-trained model within a transfer learning framework, utilising experimentally measured absorption data.

3.7 Conclusions

This chapter has presented the systematic methodology employed to design, fabricate, and experimentally evaluate the acoustic absorber device investigated in this thesis. The design stage utilised an analytical modelling framework to provide an efficient mathematical description of the absorber. This will enable rapid and cost-effective data acquisition, as well as fast preliminary characterisation to support the investigations presented in the subsequent chapters. Fabrication was carried out using FDM technology with PLA material. Subtle variations in key printer parameters, such as nozzle temperature, layer height, and infill density, were made to improve the quality of manufactured parts. The adoption of FDM offered several advantages: rapid prototyping of intricate absorber geometries, cost-effective production, and flexibility for iterative refinement during the design-validation cycle.

Experimental testing procedure using a two-microphone impedance tube setup, was described, enabling precise measurement of normal-incidence absorption spectra over the target frequency band. To ensure reliability, special emphasis was placed on experimental reproducibility, accurate calibration of the measurement system.

A DL framework, referred to as Structured Latent Autoencoder (SLAE), was introduced to establish a bidirectional mapping between geometry and acoustic performance. In the inverse mode, the SLAE predicts geometric parameters from measured or simulated spectral properties. In the forward module, it reconstructs the input spectrum from its latent (geometry) representation. To improve computational efficiency without sacrificing prediction accuracy, the original 901D analytical spectra were post-processed to extract a reduced set of key spectral features, thereby lowering input dimensionality while retaining the key physical information.

To bridge the gap between idealised analytical data and real-world measurements, physics-informed data augmentation strategies were presented. These augmentations aim to mimic plausible deviations in spectral characteristics arising from manufacturing tolerances and unmodeled physical effects. This strategy improves the network's ability to generalise beyond the purely analytical domain. Subsequently, a TL procedure was proposed to incorporate real-world manufacturing artefacts, through experimental data, into the network training. This process aims to incorporate the effects of fabrication imperfections into the model, enhancing its robustness and predictive accuracy in practical deployment scenarios.

The methodology described in this chapter provides a coherent framework for integrating analytical modelling, additive manufacturing, experimental characterisation, and DL into a unified workflow for acoustic absorber design. The combination of model-driven design, targeted fabrication, precise measurement, and intelligent data-driven modelling lays the foundation for the comprehensive results and discussions presented in the subsequent chapters.

Chapter 4

Model Validation and Parametric Studies

4.1 Introduction

In Chapter 3, the absorber model underpinning this thesis was presented in detail, encompassing its conceptual design, analytical formulation, and numerical simulation framework. Before any practical implementation or further optimisation, it is crucial to thoroughly validate and characterise the acoustic performance of the proposed HRAM absorber. Validation ensures the model's accuracy, while characterization provides critical insights into the absorber's operational effectiveness across the intended frequency range. This chapter is dedicated to the comprehensive validation and performance evaluation of the designed acoustic metamaterial absorber (AMM). The objectives of this chapter are threefold: 1) to assess the absorber's sound absorption characteristics under experimental conditions using the two-microphone impedance tube method, 2) to verify the accuracy of the analytical and finite element models by comparing their predictions with experimental measurements, and 3) to identify and analyse any discrepancies between the modeled and measured responses, investigating their potential causes, including fabrication tolerances, material properties, and unmodeled physical effects. The validated model will serve as the foundation for the subsequent chapters, which focus on leveraging deep learning (DL) methodologies to optimize AMM performance, improving robustness against manufacturing uncertainties and enhancing the model's predictive accuracy. The results presented in this chapter are largely drawn from the author's peer-reviewed publication [177].

4.2 Numerical modelling

The HRAM design, inspired by many designs found in the literature and presented in Fig. 3.2, is modelled using finite element analysis in COMSOL[®] multiphysics. Two numerical model setups are considered here, as shown in Fig. 4.1.

Fig. 4.1(a) models the case of a single unit with a cylindrical manifold defined at the top of the absorber, having the same size and cross-section as the experimental test tube. Fig. 4.1(b) models the case where multiple absorber units are arranged in parallel, with the manifold having the same size and cross-section as each absorber unit, and each unit covers approximately $1/n$ of the tube's cross-section, where n is the number of units. It is noteworthy that the designed absorber has a hexagonal structure, while the impedance tube is cylindrical, thus, the cross-section of the material would be smaller than that of the tube, resulting in a lower experimental absorptivity compared to the actual absorptivity of the




Figure 4.1: Numerical model setups for (a) a single unit of the absorber with a cylindrical inlet manifold of the same size and cross-section as the experimental tube, and (b) multiple homogeneous units with a hexagonal inlet manifold of the same size and cross-section as each absorber unit

material. Since this would not be a true reflection of the sound absorption performance of the absorber, multiple units are printed on a panel to span the incident cross-section with each unit spanning approximately $1/n$ of the tube's cross-section, which is the case modeled by the setup of Fig. 4.1(b).

4.2.1 Test case 1

Consider a geometric sample of dimensions $L_s = 35$ [mm], $t = 0.8$ [mm], $p = 2.7$ [mm], $H = 20$ [mm], $tp = 1$ [mm], $r_o = 2.3$ [mm], and $Nt = 3$, is designed in CAD and additively manufactured using polylactic acid (PLA) on a Prusa mini machine. Fig. 4.2 shows the printing process and the test sample. A single unit of the absorber is fabricated on a 127 [mm] cylindrical panel, and tested in the impedance tube shown in Fig. 4.2(c). Note that the system has been printed as a single piece to mitigate the problem of acoustic leakage from glued parts. This is facilitated by the acoustic inlet (perforation) having the same geometry as the first turn of the spiral. Fig. 4.2(d) compares the experimental absorption curve to that of numerical based on the setup in Fig. 4.1(a). It can be observed that the absorption curves match sufficiently in this case.




Figure 4.2: (a) The 3D printing process used to fabricate the meta-absorber, illustrating the printing setup and methodology, (b) the completed fabricated sample showcasing the final design, (c) the experimental text rig, and (d) shows the experimental result compared to the numerical absorption curve.

4.2.2 Test case 2

A second sample with parameters $L_s = 50$ [mm], $t = 0.5$ [mm], $p = 7.2$ [mm], $H = 13$ [mm], $tp = 1$ [mm], $r_o = 2.5$ [mm] and $Nt = 4$, is additively manufactured to evaluate the current mesh settings at an extreme sub-wavelength frequency of 111 [Hz]. A discrepancy in the absorption amplitude between experimental and numerical results, as shown in Fig. 4.3(a), is observed. Fig. 4.3(c) depicts the current mesh setup, as detailed in Section 4.2. While this difference could typically be attributed to manufacturing imperfections or tolerances, the velocity and energy dissipation (Fig. 4.4) provide crucial information. Figure 4.4(a) shows the instantaneous total acoustic particle velocity magnitude, revealing strong flow acceleration at the aperture entrance. This behaviour reflects the resonant nature of the absorber, whereby acoustic energy is effectively trapped and funneled within the internal cavity-channel system. This elevated velocity indicates regions of strong acoustic motion and energy con-

centration. Figure 4.4(b) shows the total thermoviscous power dissipation density, showing that acoustic energy loss is predominantly localised along the channel walls rather than in regions of maximum velocity. This confirms that absorption at resonance is governed by thermoviscous loss mechanisms arising from viscous shear and thermal conduction within the acoustic boundary layers adjacent to the solid surfaces. When the boundary layer thickness becomes comparable to the characteristic channel dimensions, these losses are significantly enhanced. To this end, the mesh density at the inlet and within the channels may be refined for improved accuracy. Figure 4.3(b) and (d) demonstrate the effect of mesh refinement at the inlet only, with a maximum element size of 0.2×10^{-3} [m] applied to this critical region. Thus, at extreme sub-wavelength scales, the baseline mesh settings may be insufficient for accurate predictions. However, increasing the mesh density to achieve the required accuracy significantly increases the computational cost. To address this trade-off between accuracy and computational efficiency, an equivalent model of the absorber is developed in the next section.




Figure 4.3: Experimental and numerical absorption curves of the absorber with parameters $L_s = 50$ [mm], $t = 0.5$ [mm], $p = 7.2$ [mm], $H = 13$ [mm], $tp = 1$ [mm], $r_o = 2.5$ [mm] and $Nt = 4$.




Figure 4.4: Acoustic field distribution of the HRAM at resonance. (a) Instantaneous total acoustic particle velocity magnitude, showing strong flow acceleration at the aperture entrance. (b) Total thermoviscous power dissipation density, indicating that acoustic energy dissipation is concentrated along the channel walls where viscous and thermal boundary-layer effects dominate. The contrast between velocity and dissipation fields highlights that absorption at resonance is governed by thermoviscous losses near solid boundaries.

4.3 Simplified numerical model

Although numerical modeling plays a crucial role in engineering and scientific simulations, enabling detailed analysis of complex systems, a high-fidelity numerical models, such as those based on FEA, often come with significant computational costs. These models require extensive processing power and memory, making real-time simulations and iterative design optimization challenging. One approach to alleviate this computational burden is to approximate complex structures with simplified representations that facilitate efficient meshing and reduced computation while preserving the essential behavior of the original system. Thus, to reduce the complexity involved with the FW absorber system, a simplified numerical model is developed for improved efficiency. Here, the spiral is modeled as a straight rectangular duct connected to an equivalent cylindrical cavity with an equivalent radius of $r_{eq} = \sqrt{(1.5\sqrt{3}Ls^2)/\pi}$. The cavity is determined such that the volume occupied by the spiral is accounted for. To this end, the effective depth of the cavity is the depth, H , scaled by the volume fraction, v_f , where,

$$v_f = \frac{V_c - V_{sp}}{V_c}, \quad (4.1)$$

$V_c = S_2 H$ is the volume of the cell, S_2 is the cross-sectional area of the hexagonal cell, and V_{sp} is the volume of the spiral. The length of the straight duct is calculated along the effective flow path of the coiled channel as shown in Fig. 4.5(b), so that the total length of the duct is calculated as $(L_2 + tp + \epsilon)$, where L_2 is the length along the effective flow path, and ϵ accounts for corrections at both ends of the channel due to geometric discontinuities. Thus, for an arbitrary aperture, a simple correction $\epsilon = 0.96\sqrt{S_1}$ [158] is used, where $S_1 = pH$, is the cross-sectional area of the channel. The expressions for V_{sp} and L_2 , will be presented in a later section. A cylindrical manifold, highlighted in blue in Fig. 4.5(a), with the same cross-sectional area as the cavity is defined above the duct. Thermo-viscous physics is applied to both the duct and cavity, while pressure acoustics is applied to the manifold. The system is then meshed following the same procedure described in Section 4.2 and a unit pressure excitation is applied to the inlet manifold in the -z direction. To further reduce the computational load, a quarter section of the system is modeled, leveraging symmetry, as shown in Fig. 4.5(a). Fig. 4.6 presents the acoustic absorption curves comparing experimental result, FW simulation, and the SNM. Simulations are performed based on the setup shown in Fig. 4.1(b). The cavity side length of $L_s = 35$ [mm] allowed four (4) cells to be manufactured on a 199.5×199.5 [mm] square panel for testing. The SNM closely matches the FW simulation while significantly reducing computational time to approximately 80 minutes, compared to approximately 6 hours required for the FW simulation, and around 96 hours after mesh refinement at the inlet. The SNM thus offers a valuable and efficient tool for low-cost optimization studies.

4.3.1 Effect of the number of absorber units on acoustic response

To investigate the effect of the number units on acoustic response, four square panels of size 199.5×199.5 [mm], each having different number of homogeneous units (schematised in Fig. 4.7) are additively manufactured. A maximum of four units can be printed on the square panel due to experimental constraints. Fig. 4.8 show the results. Note that the setup of Fig. 4.1(b) has been used for the multi-cell case in the numerical simulation. It is observed that the peak absorptivity in the experimental measurements increases as more units are printed on the panel. This increase continues until a maximum peak absorptivity is reached, which aligns with the numerical prediction of an infinite surface. The convergence of numerical and experimental absorptivities at four units in this case suggests that this




Figure 4.5: Equivalent model of the absorber showing: (a) a quarter-section view of the equivalent model comprising a cavity and an equivalent straight channel which includes the perforation. The manifold is highlighted in blue, and (b) a schematic representation of the coiled channel, with the effective length calculation path explicitly marked to demonstrate the acoustic pathway.

configuration effectively matches the simulated infinite surface and beyond this point, edge effects are negligible, and the acoustic response becomes stable.




Figure 4.6: Comparison of absorption performance curves obtained from the simplified numerical model (SNM), full-wave (FW) simulations, and experimental measurements. The absorber parameters used in the simulations are $L_s = 35$ [mm], $t = 0.8$ [mm], $p = 2.7$ [mm], $H = 20$ [mm], $r_o = 2.3$ [mm], and $N_t = 3$. The figure highlights the alignment between the simplified model and the full-wave simulations, as well as their agreement with experimental results, demonstrating the validity and accuracy of the equivalent model for predicting absorber performance.




Figure 4.7: Schematic arrangement of unit cells on 199.5 [mm] by 199.5 [mm] square panels. These are additively manufactured and responses are compared and shown in Fig. 4.8.




Figure 4.8: Comparison of experimental absorption curves for units 1 through 4 (shown as circular markers) and the numerical absorption curve (depicted by the black solid line). The plots show that absorptivity increases with increase in the number of units printed on the panel and convergence is reached at infinity, corresponding to the numerically stimulated result with infinite surface assumption. Thus, four units are optimal and infinite in this case to realize the true absorption performance of the absorber.

4.4 Experimental validation

The impedance tube measurements are performed according to ISO 10534-2 standards, as detailed in Section 3.5. Sample geometries are simulated and additively manufactured (see Table 4.1 for the geometric values), and the air parameters used in the models (numerical and analytical) are presented in Table 4.2. Fig. 4.9(a) shows that the analytical result matches well with those of numerical and experiment in terms of the peak absorptivity and resonant frequency. The 199.5×199.5 [mm] panel allowed for four (4) homogeneous absorber units, each of cavity size $L_s = 35$ [mm], to be fabricated on it, taking into consideration the tube's circular cross-section. The bandwidth of the analytical curve is noticeably narrower and slightly overestimates the absorptivity compared to the numerical and experimental results. This discrepancy may be attributed to unmodeled physics in the absorber system, arising from complex wave interactions between the component resonant structures, i.e., the cavity and the coiled channel. As previously described and shown in Fig. 4.4, the perforation dominates the damping as the particle velocity changes more rapidly there, especially around the sharp curvature, while in the rest part of the channel, no significant friction loss is observed as the air particle moves comparably slower. Thus, the coiled channel functions as the neck of a Helmholtz resonator where the sound energy is consumed by fierce flow friction. The second validation case in Table 4.1 aims to achieve a compact design in both lateral and vertical dimensions. Fig. 4.9(b) shows the result which achieves a near perfect absorption at 204 [Hz], with a lateral side length, L_s , of 30 [mm] and a thickness of 15 [mm]. However, it can be observed that both the numerical and analytical bandwidths fail to match the experimental results. More so, the numerical curve slightly underestimates the absorptivity, despite mesh refinements. These deviations could be due to over-simplified modeling assumptions for the current system and/or artifacts introduced during the manufacturing process. The third validation case aims to achieve a near-perfect absorption at a much lower resonant frequency and sub-wavelength thickness. Here, only three units can be manufactured on the panel as shown in Fig. 4.9(c) due to experimental constraints. Again, the narrower bandwidths observed in the analytical and numerical absorption curves persist. These deviations are primarily attributed to the limitations of traditional acoustic modeling, which fails to capture the complexities of the current system's behavior, and the impact of additive manufacturing on the fabricated sample is also a contributing factor, as the spiral wall thickness $t = 0.5$ [mm] used in this case is sufficiently thin, which may compromise

its acoustic rigidity as assumed in simulation. This discrepancy likely introduces additional acoustic damping and peak broadening due to structural excitation, thereby contributing to the observed differences. The higher absorptivity observed in the analytical curve is largely

Table 4.1: Sample geometric parameters of the simulated absorber with $t_p = 1$ [mm] in all cases

Sample no	p	t	Nt	Ls	H
1	2.7 [mm]	0.8 [mm]	3	35 [mm]	20 [mm]
2	2.2 [mm]	0.8 [mm]	2.5	30 [mm]	15 [mm]
3	7.2 [mm]	0.5 [mm]	4	50 [mm]	16 [mm]
4	6.7 [mm]	0.8 [mm]	4	45 [mm]	18 [mm]

due to over-estimation. Nonetheless, a near-perfect absorption is achieved at 111 [Hz] with an extremely sub-wavelength thickness of about $1/193$ of the operating wavelength ($H = 16$ [mm]). Notably, the absorber's performance is not governed by classical quarter-wavelength theory but rather by a hybrid resonance mechanism combining the coiled channel (spiral) and Helmholtz cavity effects. Consequently, the effective path length of the wave inside the coiled channel is approximately $\lambda/9$, which is significantly shorter than that required by conventional quarter-wavelength resonators. The fourth case also shows significant disparities in the absorption bandwidths. These results shows the inconsistencies in performance across different simulated geometries. It is important to note that models (numerical or analytical) are simplified representations of real systems and are often sufficient within a chosen operational range. However, the behavior of the absorber system underscores the need for advanced modelling methods that ensure robustness and consistency across the entire operational range for precise noise control.

Table 4.2: Air simulation parameters used in all models

Symbol	Name	Value	Unit
η	Dynamic viscosity	1.825×10^{-5}	Pa.s
ρ	Density of air	1.204	kg/m^3
c	Speed of sound in air	343	m/s
γ	Specif heat ratio	1.4	-
κ	Thermal conductivity	0.0257	W/(mK)
Cp	Heat capacity at constant pressure	1005	J/(kgK)




Figure 4.9: (a) Comparison of the analytical absorption curve with numerical and experimental results. The insets show the CAD designs of the four geometric cases from Table 4.1 that are fabricated.

4.5 General design considerations

A design guideline to achieving a suitable trade-off between absorber thickness and lateral size for a compact device, is provided in this section. Figs. 4.10(a) and 4.10(b) present the contour plots illustrating the resonance behavior for varying thickness and cavity size. Keeping other parameters constant, larger values of H and L_s tend to yield higher peak absorptivity, and lower resonance frequency. This phenomenon occurs because larger H allows for increased interaction (or exposed surface area) between the sound wave and the absorber, and longer acoustic wavelengths (associated with lower frequencies) are much favoured with larger cavities. The color map represents the magnitude of the plotted variables, with different colors indicating different levels of H and L_s . Each color strip connects approximately same levels of H and L_s . Thus, an absorber with a target frequency of 180 [Hz] and peak absorptivity of ≈ 1 would have L_s of about 35 [mm], and a thickness of about 17 [mm]. Similarly, the influence of the number of turns, N_t , on the resonance behavior is shown in Fig. 4.10(c).

Increasing N_t lowers the resonance frequency until $N_t = 6$, primarily due to the elongation of the channel length. However, beyond $N_t = 6$, the cavity dominates and resonance switches to higher frequencies due to a substantial reduction in effective cavity volume. Subsequently, absorption performance deteriorates and eventually plateaus. In fact, for design purposes, $2 \leq N_t \leq 4$, is sufficient.




Figure 4.10: Influence of size-related parameters and the number of turns on resonance behavior. (a) and (b) are contour plots illustrating the resonance behavior of the absorber at varying values of L_s and H , while keeping other geometric parameters constant: $t_p = 1.5$ [mm], $N_t = 3$, $p = 5$ [mm], and $t = 0.8$ [mm], and (c) illustrates the effect of varying N_t on resonance behavior, with the parameters $t_p = 1.5$ [mm], $p = 4.2$ [mm], $t = 0.8$ [mm], $L_s = 63.5$ [mm], and $H = 15$ [mm], held constant. This highlights the interplay between L_s and H as well as the impact of increasing the channel length through N_t on the resonant frequency.

4.5.1 Sensitivity analysis

The effects of channel width, p , and wall thickness, t , on absorption peak are investigated, as shown in Fig. 4.12. The parameter, p , is varied in the range of 5.7 - 7.2 [mm] in increment of 0.5 [mm], and t is varied in the range of 0.8 - 2.5 [mm] in increment of 0.5 [mm] so that the rates of increase in p and t are 7.7 % and 62.5 %, respectively. During the simulation,

parameters $Nt=3$, $Ls=50$ [mm], $H=20$ [mm], and $tp=1.5$ [mm], are fixed as defined. It is clear that the shift in resonance frequency is more sensitive to variations in t than in p , and decreasing (increasing) t moves the peak to higher (lower) frequency, while decreasing (increasing) p moves the peak to lower (higher) frequencies, as shown in Figs 4.11(a) and 4.11(b). This shows that the spiral wall thickness is significant in controlling the acoustic behaviour of the absorber. The impedance matching condition for perfect absorption is realised when $\text{Re}(Z/Z_o) = 1$ and $\text{Im}(Z/Z_o) = 0$. It can be observed from Fig. 4.12 that the imaginary parts cross zero, a signature of resonance, however, the real parts are not equal to Z_o in most cases, implying partial reflection and imperfect absorption (less than unity)




Figure 4.11: Absorber responses to variations in (a) p and (b) t . The absorber responses are more sensitive to changes in t than in p , implying that the wall thickness plays a significant role in shaping the acoustic response of the absorber.




Figure 4.12: The real and imaginary parts of the impedance for different channel widths, p , and spiral wall thickness, t . Zero crossings of the imaginary part indicate resonances, though absorption remains below unity in most cases due to incomplete impedance matching

4.6 Conclusions

This chapter has presented a comprehensive, multi-method validation of the proposed AMM absorber, integrating analytical, numerical, and experimental approaches. The primary aim was to verify the reliability of the design's predicted performance while identifying and understanding discrepancies arising from modelling assumptions, computational limitations, and fabrication realities.

The FEM simulations provided a high-fidelity representation of the absorber's acoustic behaviour, capturing complex physical phenomena such as viscothermal losses, higher-order mode interactions, and non-uniform pressure fields, effects that may be neglected in analytical models. Numerical results provided good qualitative agreement with analytical predictions, particularly in the location and magnitude of peak absorption. Minor deviations, such as small shifts in resonance frequency and modest variations in absorption bandwidth, were

traced to geometric features and dissipative mechanisms not accounted for in the analytical framework. Given the high computational cost of full-wave (FW) FEM simulations, approximately 6 hours per configuration, a simplified numerical model (SNM) was developed to reduce computational demand. This reduced-order model achieved comparable accuracy, while lowering the runtime to around 80 minutes. However, despite this improvement, the computational cost remained prohibitive for rapid parametric studies or downstream applications such as DL-based optimisation, which require the generation of large datasets in a short time.

The analytical model offered a computationally efficient alternative for predicting the absorber's performance. Its mathematical simplicity allowed for clear insight into the influence of geometric parameters on the overall absorption behaviour, making it particularly well-suited for sensitivity analyses and preliminary design screening. Nevertheless, the assumptions inherent in the TMF, combined with unmodelled wave phenomena within the absorber, resulted in measurable discrepancies when compared with high-fidelity FEM results and experimental measurements.

The experimental validation, carried out using impedance tube measurements in accordance with ISO 10534-2, confirmed the overall performance trends predicted by both models. While the peak absorption values and general spectral profiles were consistent, small discrepancies in absorption magnitude and resonance position were observed. These deviations were attributed, on one hand, to manufacturing imperfections that were not represented in either the numerical or analytical models, and, on the other hand, to unmodelled physical phenomena in the theoretical and computational frameworks.

Across all three validation approaches, the absorber demonstrated strong sub-wavelength performance, confirming the feasibility of the proposed design for targeted sound attenuation. The alignments between analytical, numerical simulation, and experimental measurement provide sufficient evidence that the basic underlying physics of the absorber has been captured. Importantly, the identified sources of discrepancy highlight areas where design robustness must be addressed—particularly in accounting for manufacturing imperfections and variability in real-world conditions.

The findings from this chapter form a critical foundation for subsequent chapters. The established validation framework not only confirms the performance of the baseline AMM design but also provides the necessary insight into its sensitivity to parameter variation.

This understanding is essential for the development of the DL-based strategies implemented in subsequent chapters, where the goal is to optimise performance while enhancing robustness against fabrication uncertainties.

Chapter 5

Traditional and Deep Learning Methods in AMM Design

5.1 Introduction

The optimisation of acoustic metamaterial (AMM) absorbers remains a central challenge in the pursuit of compact, efficient, and tunable noise control solutions. While traditional optimisation techniques, such as gradient-free search algorithms, genetic algorithms, and gradient-based routines integrated into commercial solvers, have been widely applied in conventional acoustic absorber design for several decades, their use in AMM research is more recent, coinciding with the emergence of metamaterials in the early 2000s. These approaches have been employed to tune geometric parameters and resonant features to achieve desired absorption characteristics.

A central consideration in the use of DL for AMM design lies in the choice of representation for the acoustic response. Traditional optimisation methods often rely on metrics derived from the absorption spectrum, such as peak absorption, resonance frequency, or bandwidth. In contrast, DL models can learn from either full spectral data or condensed information depending on the desired balance between model accuracy, computational efficiency, and training data requirements. Determining the most effective input representation remains an important methodological question, with implications for both prediction quality and optimisation speed.

This chapter provides a comparative study of traditional optimisation techniques and DL-based approaches for AMM absorber design. The comparison will be made across several key metrics, including accuracy, time complexity, scalability, and robustness against uncertainties. The objective is to elucidate the relative strengths and limitations of each approach, and to highlight scenarios in which one methodology may be more advantageous than the other. Ultimately, this analysis will provide a framework for selecting the most appropriate optimisation strategy depending on the application context, computational resources, and robustness requirements of the AMM design process. The content of this chapter is largely based on prior work in the conference [167] and subsequent studies [190]. All simulations in this chapter were performed in MATLAB on a workstation equipped with an Intel® Xeon® W-2265 CPU @ 3.50 GHz (12 cores, 24 threads) and 256 GB of RAM.

This chapter is structured as follows. Section 5.2 introduces the design principles of AMMs and presents a detailed formulation of the generalized pattern search (GPS) traditional optimisation method, along with the deep learning architecture employed in this study. Section 5.4 presents a comparative analysis of the two approaches, using robustness and computa-

tional time as evaluation metrics. Section 5.5 examines the impact of different input data configurations on the predictive performance of the deep learning model. Finally, the main conclusions of the chapter are summarized in Section 5.6.

5.2 Design principle

The absorber developed in this thesis is guided by the principle that acoustic metamaterials derive their functionality not from the intrinsic properties of their constituent materials, but from carefully engineered sub-wavelength structures. By exploiting tailored resonances, the design enables strong control of wave–structure interactions within a compact form factor, while remaining amenable to fabrication by additive manufacturing. A central challenge lies in balancing acoustic performance with manufacturability. The design space of AMMs is inherently high-dimensional, with performance dictated by the delicate interplay of geometric parameters, resonant mechanisms, and viscothermal losses. Exploring this space exhaustively using finite element analysis is computationally prohibitive, while fabrication imposes additional constraints related to feature resolution, robustness to imperfections, and scalability of production. These considerations highlight the need for absorber concepts that combine theoretical efficacy with structural simplicity, robustness against uncertainty, and suitability for experimental demonstration. In this thesis, the proposed HRAM absorber design reflects these requirements: it is geometrically rich enough to capture the complex resonance phenomena that challenge conventional modelling, yet structurally simple to enable optimisation, experimental validation, and integration into data-driven design workflows. This pragmatic design principle ensures that subsequent investigations—numerical, analytical, experimental, and deep-learning-based—are grounded in a model system that is both scientifically tractable and practically realisable.

5.3 Traditional optimisation

Traditional optimisation methods can broadly be categorised into three groups: gradient-based algorithms, gradient-free local search methods, and global optimisation techniques. Gradient-based methods, such as steepest descent and quasi-Newton routines, rely on the evaluation of derivatives of the cost function with respect to design parameters. These approaches are computationally efficient when the optimisation landscape is smooth and differentiable, but they can easily become trapped in local minima, an important limitation

in AMM design where the response surfaces are often highly nonlinear and multimodal. Gradient-free local search methods, such as the Nelder–Mead simplex algorithm, overcome this reliance on derivatives and can operate on noisy or discontinuous functions, making them attractive for FEM-based optimisation. However, their local nature restricts their ability to explore large design spaces effectively. Global optimisers, on the other hand, such as genetic algorithms (GA), particle swarm optimisation (PSO), and generalized pattern search (GPS), provide a more systematic exploration of the design space and are better suited for handling complex, multimodal optimisation problems characteristic of AMM design. Among these, GPS is particularly noteworthy because it combines the robustness of global search with the simplicity of derivative-free evaluation. GPS proceeds by exploring candidate solutions on a structured mesh that is iteratively refined, allowing for efficient convergence without the need for gradient information. Its deterministic nature makes it more predictable than stochastic global optimisers such as GA or PSO, while still retaining robustness to noise and non-linearity. Given these attributes, GPS is selected in this study as the representative traditional optimisation method for comparison against deep learning-based approaches, particularly in terms of accuracy, computational cost, and robustness to modelling and fabrication uncertainties.

5.3.1 Generalised pattern search

Numerous traditional optimisation algorithms exist in the literature, and to facilitate the current study, a non-stochastic GPS algorithm, proposed by Torczon [191], is deployed. Although genetic algorithm (GA) is more commonly deployed for AMM design [192, 193], a different approach is adopted in this study, which is comparatively simpler to implement and avoids the need for defining genetic operators, parameter tuning, and managing population dynamics. GPS is a global derivative-free optimisation algorithm that explores a set of points, called a mesh, around the current point at each iteration. This mesh is constructed by adding the current point to a scalar multiple of vectors, called patterns. GPS polls (computes the objective function) the points in the current mesh, and the exploration is deemed successful if the function value at any point is smaller than the previous one. One GPS version of Torczon with a maximal positive basis of $2N$ vectors is adopted here, where N is the number of decision variables. In this case, the mesh points for $N=4$ variables are defined, and the eight pattern vectors are each scaled by Δ . The mesh size is controlled by Δ , and is increased by a factor of 2 for every successful poll, or halved if the poll is unsuccessful. The search

begins by initialising p , Nt , Ls and H , with $\Delta_0 = 1$. Here, the target frequency, f_o and $t = 0.8$ [mm], are fixed for simplicity, while the algorithm searches for the optimal output, $\mathbf{y}$, as it maximises α at a given target frequency. r_o is determined based on the values of p and t by solving Eq. (3.1) for $r_o > 0$.

Consider the optimisation of the absorber parameterised by four design variables, collected into the vector

$$\mathbf{u} = \begin{bmatrix} p & Nt & Ls \end{bmatrix}^T \in \mathbb{R}^3, \quad (5.1)$$

the optimisation problem is posed as

$$\min_{\mathbf{u} \in \Omega} J(\mathbf{u}), \quad (5.2)$$

where $J(\mathbf{u}) = \frac{1}{\alpha}$ (α is the absorptivity defined in Eq. (3.9)) is the objective function defined at a target frequency value. The feasible set, Ω , is defined by bound constraints on each design variable:

$$\Omega = \{ \mathbf{u} \in \mathbb{R}^3 \mid u_i^{\min} \leq u_i \leq u_i^{\max}, \quad i = 1, \dots, 3 \}. \quad (5.3)$$

Within the GPS framework, let $\mathbf{u}_k \in \mathbb{R}^3$ denote the current iterate and $\Delta_k > 0$ the current mesh size. The polling directions are defined as the maximal positive basis in three dimensions:

$$\mathcal{D} = \begin{bmatrix} I_3 \\ -I_3 \end{bmatrix} \in \mathbb{R}^{6 \times 3}, \quad (5.4)$$

which corresponds to the eight pattern vectors $\{\pm e_1, \pm e_2, \pm e_3\}$. The corresponding poll set at iteration k is

$$\mathcal{P}_k = \{ \mathbf{u}_k + \Delta_k \mathbf{d}_j \mid j = 1, \dots, 6 \}, \quad (5.5)$$

or equivalently, in matrix form,

$$M_p = \mathbf{1}_6 \mathbf{u}_k^T + \Delta_k \mathcal{D} \in \mathbb{R}^{6 \times 3}, \quad (5.6)$$

where $\mathbf{1}_6 \in \mathbb{R}^6$ is an all-ones column vector and the j -th row of M_p corresponds to the poll point

$$\mathbf{p}_j = \mathbf{u}_k + \Delta_k \mathbf{d}_j. \quad (5.7)$$

For a broadband response, a supercell composed of n inhomogeneous units is defined by the set $\mathbf{U} = \{ \mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n \}$, and the objective function is posed to maximise the average absorptivity over a given frequency range. The average absorptivity over the frequency band

of interest $\mathcal{F} = [f_{\min}, f_{\max}]$, is

$$\bar{\alpha}(\mathbf{U}) = \frac{1}{|\mathcal{F}|} \int_{\mathcal{F}} \alpha(\mathbf{U}, f) \, df, \quad |\mathcal{F}| = f_{\max} - f_{\min}, \quad (5.8)$$

where $\alpha(\mathbf{U}, f) \in [0, 1]$ denotes the absorptivity of the supercell $\mathbf{U}$ at frequency f . To simultaneously maximise the average absorption and minimise the variation across the band (i.e., to keep resonances close together or avoid deep dips between peaks), a composite objective function may be defined to minimise

$$J(\mathbf{U}) = -\bar{\alpha}(\mathbf{U}) + \Theta \sigma_{\alpha}(\mathbf{U}), \quad (5.9)$$

where $\Theta > 0$ is a weighting factor controlling the trade-off between high mean absorption and flat broadband performance. Here, the discrete average absorption over the frequency band is given by

$$\bar{\alpha}(\mathbf{U}) = \frac{1}{m} \sum_{i=1}^m \alpha(\mathbf{U}, f_i), \quad (5.10)$$

and the standard deviation of absorption, quantifying the spread of the response, is

$$\sigma_{\alpha}(\mathbf{U}) = \sqrt{\frac{1}{m} \sum_{i=1}^m \left(\alpha(\mathbf{U}, f_i) - \bar{\alpha}(\mathbf{U}) \right)^2}, \quad (5.11)$$

where m is the length of the frequency grid over which the mean absorptivity is maximised.

Rationale and implications of deterministic global optimisation

The GPS algorithm was selected as the representative traditional optimisation method in this study due to its deterministic, derivative-free nature and its suitability for highly non-linear and multimodal design spaces typical of acoustic metamaterials. Unlike stochastic global optimisers such as genetic algorithms or particle swarm optimisation, GPS produces reproducible solutions for a given problem formulation, mesh definition, and initialisation. This determinism was deliberately chosen to ensure consistent benchmarking against deep learning models and to isolate performance differences attributable to methodology rather than random sampling effects. However, it is important to recognise that determinism does not imply global optimality nor independence from process initialisation. The convergence behaviour of GPS is influenced by the choice of the initial design point, the initial mesh size and the mesh refinement strategy. Different initialisations may lead the algorithm to distinct

locally optimal basins within the feasible design space, particularly in problems exhibiting strong non-convexity and multiple resonant mechanisms, as is the case in AMM design.

In this context, the decency of a solution refers not solely to the attainment of a high objective value, but to the stability, reproducibility, and physical plausibility of the resulting design under realistic perturbations. While GPS often yields solutions with near-optimal peak absorptivity, these solutions may exhibit heightened sensitivity to parameter variations, reflecting convergence to narrow optima. This behaviour motivates the subsequent robustness analysis and highlights an inherent trade-off between peak performance and tolerance to uncertainty.

The deterministic nature of GPS therefore provides a controlled and interpretable optimisation baseline, while also exposing limitations related to initialisation dependence and local convergence. These characteristics make GPS a suitable but deliberately constrained comparator for assessing the advantages of data-driven design approaches, particularly in terms of robustness, generalisation, and scalability.

5.3.2 Deep learning model

The training data for the subsequent comparative study are generated using the analytical model of the meta-absorber. Geometric parameters are sampled from a uniform distribution within the ranges specified in Table 3.3. A fully connected deep neural network (DNN) is then trained on this dataset. The network architecture consists of four fully connected layers with the following number of units: $3 \rightarrow 100 \rightarrow 50 \rightarrow 25 \rightarrow 3$. All hidden layers use the ReLU activation function, while the output layer employs a linear activation. The input to the network is $\mathbf{u} = [\alpha, f_r, H]^T$. H is included in the input vector to address the non-uniqueness problem of inverse mapping. The network output is $\mathbf{y} = [p, Nt, Ls]^T$. Data were collected from the analytical model described in Section 3.2, and are divided as described in Fig. 3.6. The network was trained using the *scaled conjugate gradient (SCG) algorithm* [194]. SCG is an efficient second-order optimization method that approximates the Hessian without performing computationally expensive line searches, making it suitable for large datasets and deeper networks. The weight update in SCG is given by:

$$w_{k+1} = w_k + \alpha_k d_k \quad (5.12)$$

where w_k is the weight vector at iteration k , d_k is the conjugate search direction, and α_k is the step size determined by the scaling procedure.

The search direction is updated according to:

$$d_{k+1} = -g_{k+1} + \beta_k d_k \quad (5.13)$$

where $g_{k+1} = \nabla E(w_{k+1})$ is the gradient of the loss function E with respect to the weights, and β_k is a scaling factor computed from the local curvature approximation.

The network architecture is shown in Fig. 5.1. $\mathbf{y}^1, \dots, \mathbf{y}^n$ are the outputs of layers 1 through n , $\mathbf{W}^1, \dots, \mathbf{W}^n$ and $\mathbf{b}^1, \dots, \mathbf{b}^n$ are the layer weight matrices and bias vectors, respectively. The output of a layer, L , is given by $\mathbf{y}^L = \text{ReLU}(\mathbf{W}^L \cdot \mathbf{y}^{L-1} + \mathbf{b}^L)$, $\forall L = 1, 2, \dots, n-1$, and $\mathbf{y}^0 = \mathbf{u}$. The output of the n th (last) layer is $\mathbf{y}^n = \text{Linear}(\mathbf{W}^n \cdot \mathbf{y}^{n-1} + \mathbf{b}^n)$, where the activation "Linear(.)" is an identity. The network training was performed in MATLAB using its Deep Learning Toolbox.




Figure 5.1: Architecture of the DNN with layer input and output dimensions and those of the weights and biases. $\mathbf{S}^L$ denotes the number of artificial neurons in layer L , $\forall L = 1, 2, \dots, n$

Figure 5.2 illustrates the learning curve of the network. Both training and validation losses decrease sharply during the initial 500 epochs, indicating rapid learning. Between epochs 500 and 1000, the rate of decrease slows, and the curves begin to converge, suggesting the model is approaching optimal performance. After epoch 1000, the training loss stabilizes at a low level, with the validation loss completely superimposed on it. The lack of divergence between the two curves indicates no signs of overfitting, and the model maintains good generalization throughout training. Table 5.1 presents the quantitative assessment of the network performance on the three predicted outputs. It summarises the predictive performance of the DNN model on the three geometric parameters. The performance indicates that the model achieves good agreement with the reference data for parameter p ($R^2 = 82.1\%$), while prediction accuracy is moderate for Nt ($R^2 = 58.3\%$) and Ls ($R^2 = 66.0\%$). This demonstrates that the network is capable of reliably capturing underlying trends for some geometric




Figure 5.2: Training and validation learning curves of the DNN model showing rapid decrease in the loss and close convergence of the training and validation loss, indicating no signs of overfitting

features, although others exhibit greater variability or nonlinear dependencies that are more difficult to learn from the available data. Despite these limitations, the achieved accuracy remains valuable for reducing the design search space, as the model provides sufficiently accurate estimates for use in subsequent comparative studies. Several network configurations, training algorithms, and pre-processing strategies were explored, yet the performance consistently converged to a similar range, suggesting that the observed errors arise primarily from the intrinsic complexity of the problem rather than model underfitting. These results therefore establish a practical baseline for geometry–response prediction and are further enhanced in a later chapter using an inverse–forward autoencoder architecture, which facilitates data augmentation and transfer learning for subsequent studies.

Table 5.1: Quantitative assessment of DNN performance on the predicted outputs.

Output	MSE	R^2 (%)
p	8.01×10^{-3}	82.1
Nt	3.30×10^{-2}	58.3
Ls	2.20×10^{-2}	66.0

5.4 Comparative study

The absorptivity curves based on the predicted geometric parameters from DNN and GPS are compared relative to arbitrarily chosen target performances. The primary objective here is to assess the effectiveness of DL in the design of AMMs, and to determine whether the advantages of DL approaches outweigh those of traditional optimisation (TO) techniques. Alternatively, the analysis may reveal specific trade-offs that must be considered depending on the problem at hand. The comparison is based on the following key metrics: robustness, inference time, and time complexity. For validation purposes, two prediction scenarios were obtained from both the DNN and TO. Table 5.2 presents the predictions from both models based on the desired spectral and geometric targets, i.e., α_p , f_r , and H . It is noteworthy that these values were set based on some reference (experimentally measured) targets, thus, the TO optimised absorptivity may be higher. Fig. 5.3 shows the comparison plots.

Table 5.2: Comparison of predicted output values for geometric parameters from the DNN and GPS models.

Design spec 1 ($\alpha_p = 0.97$, $f_r = 110$ [Hz], $H = 16$ [mm])			
	p [mm]	Nt	Ls [mm]
DNN	4.73	3.86	43.84
GPS	6.78	4.28	46.46
Design spec 2 ($\alpha_p = 0.92$, $f_r = 151$ [Hz], $H = 20$ [mm])			
DNN	4.37	3.08	37.34
GPS	3.84	5.0	28.53

5.4.1 Robustness to parametric uncertainty

To ensure a consistent and reliable performance when the geometric parameters are subject to variations/noise, it is desired to have a robust system that is tolerant to such conditions. Such variations could arise from, for example, dimensional inaccuracies from additive manufacture, variations in environmental conditions/material properties. For fused deposition modelling (FDM) using a 0.4 [mm] nozzle and PLA filament, deviations in the final printed geometry arise from three main sources:

- **XY positional tolerance:** The printer's stepper motors allow fine positional resolution (typically ~ 0.01 [mm]), but mechanical imperfections, filament extrusion variability, and nozzle diameter effects introduce an effective uncertainty in XY dimensions, approximated as ± 0.1 [mm] for general features.




Figure 5.3: Absorption spectra obtained using the predicted geometric parameters from the DNN (black) and GPS (gray). Both predictions closely align, reproducing the targets ($f_r = 110$ [Hz] and 151 [Hz]) reasonably well.

- **Z-layer discretization:** The layer-by-layer deposition process quantizes vertical heights according to the selected layer thickness. For a 0.1 [mm] layer height used throughout this study, vertical dimensions may be effectively rounded to the nearest 0.1 [mm], with minor variations ($\sim \pm 0.05$ [mm]) due to filament compression and stepper motor precision. In this comparative study, however, H , which is related to the Z-axis movement is not predicted, but defined by the user.
- **Material shrinkage:** As PLA cools, it contracts slightly, typically 0.2 – 0.5% of the designed dimensions, introducing a multiplicative reduction across all geometric parameters.

To systematically evaluate the robustness of absorber designs against these realistic manufacturing uncertainties, a Monte Carlo (MC) simulation approach was employed. Both DL and GPS predictions were sampled under random perturbations representing the aforementioned tolerances. In each simulation, the parameters were adjusted according to:

$$p, Nt, L_s \sim \text{Nominal} \times (1 - \text{shrinkage}) + \mathcal{N}(0, \text{XY_tol}^2),$$

where $\text{XY_tol} = 0.1$ [mm], and the shrinkage factor was sampled uniformly between 0.2 – 0.5% . For each sampled geometry, the acoustic absorption was computed using the analytical model.

From the ensemble of simulated spectra, the mean and standard deviation of absorption were calculated across the frequency range of interest, providing a direct visualisation of the sensitivity of each design to manufacturing imperfections. Robustness was visualized using a $\pm 2\sigma$ prediction band around the mean, which encompasses approximately 95% of the individual perturbed samples. This approach provides a conservative estimate of variability, capturing most of the potential deviations due to manufacturing uncertainties.

Quantifying robustness

To evaluate the robustness of absorber designs under fabrication perturbations, the absorptivity within a frequency band around the target frequency $f_o = 110$ [Hz] is considered. The analytical model predicts a nominal resonance bandwidth of approximately 10 [Hz], so a conservative half-width of $\Delta f = 5$ [Hz] was used, resulting in an analysis window of $[f_o - \Delta f, f_o + \Delta f] = [105, 115]$ [Hz]. For each MC perturbed sample, the peak absorptivity within this band was extracted:

$$\alpha_{p,i} = \max_{f \in [f_o - \Delta f, f_o + \Delta f]} \alpha_i(f),$$

where $\alpha_i(f)$ is the absorption spectrum of the i -th sample. The distribution $\{\alpha_{p,i}\}_{i=1}^N$ was characterised by the mean $\bar{\alpha}_p$ and standard deviation $\sigma_{\alpha,p}$. A $\pm 2\sigma$ prediction interval,

$$[\bar{\alpha}_p - 2\sigma_{\alpha,p}, \bar{\alpha}_p + 2\sigma_{\alpha,p}],$$

was used to visualise the 95% confidence bounds of absorptivity under realistic fabrication variability. To quantify robustness relative to the absorptivity peaks, an acceptance band around the mean with $\pm 1\sigma$ was defined,

$$[\bar{\alpha}_p - \varepsilon, \bar{\alpha}_p + \varepsilon],$$

with ε being the standard deviation of perturbed samples. Two key robustness metrics were computed:

1. The fraction of perturbed samples whose peak absorptivity lies within the acceptance band around the mean, indicating stability and robustness.
2. The fraction of samples lying outside the $\pm 2\sigma$ prediction interval, providing a measure of extreme deviations due to fabrication variability.

To evaluate the robustness of each design, MC simulations with stochastic perturbations were applied to the geometric parameters. For each run, a total of 181 perturbed samples were generated, corresponding to the frequency discretisation used in simulation (0.5 [Hz] steps from 90–180 [Hz]). This procedure was repeated across 5 independent random seeds to assess statistical consistency and account for variation due to sampling noise. Thus, each design was evaluated over a total of 905 perturbed realisations (181 samples $\times$ 5 seeds). For each run, key robustness metrics were computed such as the mean peak absorptivity, standard deviation, the fraction of samples within $\pm 1\sigma$ of the mean (as a stability proxy), and the fraction of outliers beyond $\pm 2\sigma$. These metrics were then averaged across the 5 MC realisations to produce the final robustness statistics reported. Table 5.3 presents the statistical results, and Fig. 5.4 shows the visual summary.

Table 5.3: Robustness comparison between DL and GPS over 5 random seeds. Metrics are reported as mean $\pm$ standard deviation.

Metric	DL	GPS
Mean peak absorptivity	0.983 ± 0.002	0.996 ± 0.001
Standard deviation (peak)	0.034 ± 0.010	0.011 ± 0.008
Fraction within $\pm 1\sigma$ absorption band	0.976 ± 0.006	0.903 ± 0.193
Fraction outside $\pm 2\sigma$ prediction band	0.021 ± 0.006	0.008 ± 0.003




Figure 5.4: Histogram summary of robustness metrics computed over 905 MC realisations. Each bar corresponds to one robustness metric, mean peak absorptivity, peak variability, fraction within $\pm 1\sigma$, and fraction outside $\pm 2\sigma$, with error bars showing the standard deviation across runs. The visual representation allows direct comparison between DL and GPS designs across all robustness criteria.

The robustness analysis in Table 5.3 shows that the GPS design achieves a comparable variability to DL in the peak absorptivity, indicating precise and consistent performance around their optimised parameters. However, only about 90% of GPS samples fall within the $\pm 1\sigma$ absorption band, compared to the significantly higher samples ($\approx 98\%$) for DL, suggesting higher sensitivity to perturbations in certain cases for GPS. While DL has a significantly higher fraction of samples outside the $\pm 2\sigma$ prediction band compared to TO, this trade-off reflects its broader generalisation. Figure 5.5 illustrates these trends for a single MC run, showing a narrower prediction band around the mean for TO, whereas DL shows broader prediction band. To further validate the robustness analysis, the same procedure was applied to **Design spec 2** in Table 5.2, i.e., $\alpha_p = 0.92$, $f_r = 151$ [Hz], $H = 20$ [mm]), and the same trend was consistently observed.




Figure 5.5: Comparison of robustness between DL and TO designs under FDM manufacturing perturbations. Shaded regions represent $\pm 2\sigma$ prediction bands across the frequency spectrum, and solid black lines indicate mean absorption. The prediction band from GPS is slightly narrower but on the overall, both methods achieve good robustness property.

General comments

It is noteworthy that the GPS design achieves a higher peak absorptivity, which is expected since it was explicitly optimised for that objective, reaching, for example, a near-perfect absorption of 0.996. In contrast, the DL design shows a relatively lower peak value, which is not surprising given that it was not trained to fine-tune a single design point but rather to generalise over a broader distribution of geometries. Moreover, the DL absorption targets

were derived from experimentally measured values, with peak absorptivities of 0.97 and 0.92 as reported in Table 5.2. Consequently, the higher peak values obtained from GPS are not unexpected, as it was optimised without such experimental constraints.

5.4.2 Time complexity

The time it takes to make a prediction using a trained DL network is in no where comparable to that it takes to train it (training took 32.5 [min] in this case, with training data size of 250,000 samples), as predictions only involve a forward pass through the network and no parameters are learned. The inference time complexity of a DL model is primarily determined by the size and complexity of the model, as well as the size of the input data. Consider the equation for a single layer of the DNN network of Fig. 5.1, i.e., $\mathbf{y} = f(\mathbf{W}\mathbf{u} + \mathbf{b})$. The weight matrix, $\mathbf{W}$, has dimensions $S^n \times S^{n-1}$, and S^{n-1} is the number of neurons in the previous layer (input dimensionality). The input vector, $\mathbf{u}$, has size $S^{n-1} \times 1$. Thus, the time complexity of the matrix-vector multiplication, $\mathbf{W}\mathbf{X}$, is of order $O(S^n \times S^{n-1})$. The bias vector, $\mathbf{b}$, has size $S^n \times 1$. Addition of the bias vector to the matrix-vector multiplication is performed element-wise, so that the complexity is $O(S^n)$. Applying the activation $f(\cdot)$ to the result adds additional $O(S^n)$ complexity. For a single layer of the network, the overall complexity is then $O(S^n \times S^{n-1}) + O(S^n) + O(S^n)$. Thus, for a network of n layers and k input samples, the total time complexity is

$$\sum_{ns=1}^k \sum_{L=1}^n O(S^n \times S^{n-1}), \quad (5.14)$$

where, in practice, the dominant term is usually the matrix-vector operation, and other contributions are deemed relatively negligible as reflected in Eq. 5.14. Furthermore, during inference, the number of layers and neurons in each layer are constant, so that Eq. 5.14 reduces to a linear time complexity $O(k)$. Thus, for a constant sample size, k , the complexity is constant time $O(1)$. In the worst case scenario as $k \rightarrow \infty$, the computational load scales linearly with order $O(k)$.

Time complexity of the GPS

The computational complexity of the GPS algorithm is influenced by several factors, including the complexity of the objective function, the dimensionality of the search space, N , and the number of function evaluations performed per iteration. In GPS, the dominant operation is

typically the polling step, where the algorithm evaluates the objective function at a set of trial points on a mesh to identify a direction of improvement. At each iteration, GPS polls the current solution by perturbing it in a set of basis directions. Using a positive spanning set such as the $2N$ basis (which includes both positive and negative directions in each dimension), the number of polling directions per iteration is proportional to $2N$. Therefore, the computational complexity per iteration is approximately $O(N)$, assuming one objective function evaluation per direction. If the total number of iterations is T , the overall time complexity becomes $O(NT)$. However, the actual runtime also depends on the computational cost of the objective function itself which may involve expensive evaluations. As such, even with modest N and T , the total runtime can become substantial.

Procedure

The goal of this comparison is to evaluate the computational efficiency of both approaches as the input size or number of geometries increases. Each method takes an input vector $\mathbf{u}$ representing a target absorber performance specification and returns a predicted geometry. The performance of the two approaches was compared in terms of:

- **Total runtime** required to solve N inverse design cases (i.e., N geometries).
- **Average runtime per case**, computed as total time divided by N .

The following setup is used:

1. The trained DNN is used to map inputs $\mathbf{u}$ directly to absorber geometries.
2. The GPS method minimises the objective function defined in Eq 5.2 to obtain corresponding geometries.
3. For fair comparison, both methods were evaluated on identical sets of N randomly generated input vectors $\{\mathbf{u}^{(i)}\}_{i=1}^N$. The input values were sampled uniformly from physically realistic ranges: $\alpha_p \in [0.1, 1.0]$, $f_r \in [50, 500]$ [Hz], and $H \in [5, 30]$ [mm]
4. To reduce the effect of noise and variability in timing results, each experiment was repeated $M = 5$ times with different random seeds. The average runtime over $M = 5$ MC realisations is reported.

Figure 5.6 compares the total runtime of the trained DNN with the GPS algorithm as a function of input size (number of geometries).




Figure 5.6: Comparison of time complexity between TO (GPS) and (b) DNN. The computational cost of TO increases approximately linearly with input size, $O(N)$, whereas the DNN achieves a near-constant inference time, $O(1)$.

The GPS approach in Fig. 5.6(a) exhibits a clear linear trend, with runtime increasing proportionally to the number of geometries. This behavior is consistent with an $O(N)$ complexity, where each additional geometry requires a full optimisation process, resulting in cumulative growth in computational load. In contrast, the DL model in Fig. 5.6(b) achieves an empirical near-constant inference time. After an initial overhead for the first evaluation, the runtime stabilises around $7 - 9 \times 10^{-3}$ [s] approximating $O(1)$ scalability. This highlights the efficiency of the trained network, which can predict outputs for new geometries almost instantaneously, in the order of milliseconds, once the training phase is complete. Figure 5.7 shows the average time complexity per case as the number of geometries increases, providing a measure of scalability for both GPS and DL approaches. In the case of GPS (Fig. 5.7(a)), the average runtime per geometry remains nearly constant at approximately 0.31–0.32 [s]. This indicates that each new geometry requires a comparable computational effort, confirming the linear complexity observed in the total runtime analysis. Minor fluctuations are attributed to solver variability, but overall the per-case runtime does not improve significantly as more geometries are considered. In contrast, the DL model demonstrates a sharp decline in average runtime per case as the number of geometries increases (Fig. 5.7(b)). Starting from around 3×10^{-3} [s] for the first few evaluations, the average cost quickly converges to the order of 10^{-4} [s]

per geometry. This empirical trend reflects the fixed and nearly negligible inference cost of the trained network, where adding more geometries simply reduces the initial overhead. This comparison highlights a critical distinction: while GPS scales linearly and incurs a persistent per-case cost, DL inference achieves near-constant time per evaluation. The slight decrease in average runtime at small N reflects the reduction of initialisation overhead, after which the cost stabilises. This property makes DL particularly attractive for high-throughput design tasks and real-time applications, where rapid evaluation of a large number of geometries is required. Overall, while GPS provides high-fidelity solutions, its linear growth in runtime makes it computationally expensive for large-scale studies. The DL model, on the other hand, enables rapid evaluations, offering a practical trade-off between accuracy and computational efficiency, especially when large datasets or real-time applications are considered.




Figure 5.7: Time complexity plots of (a) IDDNN and (b) GPS.

Computational cost, effort, and methodological trade-offs

While the comparative results demonstrate clear differences in runtime and robustness between traditional optimisation and DL approaches, it is important to contextualise these findings in terms of both computational and human effort. Traditional optimisation methods such as the GPS require minimal upfront investment: no training data are needed, and the optimisation process can be applied directly to a new design objective with limited algorithm-

mic tuning. However, this convenience comes at a significant computational cost, as each new design query necessitates a full optimisation run, often involving hundreds or thousands of expensive model evaluations. Consequently, the time-to-solution scales linearly with the number of optimisation queries, making such methods impractical for rapid design exploration or real-time applications. In contrast, DL-based approaches incur a substantial upfront cost associated with dataset generation, network architecture selection, and training. This process requires both domain expertise and careful validation to avoid learning model bias or spurious correlations. Depending on the size of the dataset, the complexity of the network architecture, and the available computational resources, training times may range from several minutes or hours to multiple days. In the comparison performed in this chapter, the network training took about 32 [min]. Once trained, however, the marginal cost of generating new designs is negligible, with inference times on the order of milliseconds. This fundamental shift from iterative optimisation to amortised inference enables rapid design space exploration and supports scalable and interactive design workflows. From a fairness perspective, it is also important to note that DL models are not universally superior. Their performance is inherently tied to the representativeness of the training data, and they may struggle to extrapolate beyond the sampled design space. Traditional optimisation methods, while slower, do not suffer from this limitation and can adapt more readily to changes in objective functions or constraints. These observations suggest that traditional optimisation methods remain well suited for one-off design tasks where robustness and minimal setup are prioritised, whereas DL approaches are most advantageous in scenarios requiring repeated design queries, rapid evaluation, or integration into data-driven and interactive design frameworks.

Comments

The GPS-based TO used in this study exhibits a clear linear scaling of runtime with the number of geometries, as each case requires a full optimisation run. This behavior is representative of many derivative-free heuristic methods, which incur a full cost per-case. Gradient-based methods, in contrast, typically scale polynomially with problem dimension and number of iterations. When gradients are accessible, they can converge more efficiently than GPS; however, their applicability is limited in highly non-convex or noisy design spaces where local minima trap the solution. Meta-heuristic (evolutionary) approaches such as genetic algorithms or particle swarm optimisation, often have even higher computational demands, with scaling that can exceed linear growth. In comparison, the DL model achieves near-constant

inference cost per geometry, making it especially advantageous for high-throughput evaluations or real-time applications, even though its training phase is computationally intensive (about 32 [min] in this case). It is noteworthy that the MATLAB *fitnet* function used to train the DNN employs a vectorised approach to process multiple inputs simultaneously, which makes the observed inference time appear nearly constant, particularly for small to moderate batch sizes. Consequently, this empirical behavior does not reflect the theoretical worst-case linear complexity of $O(N)$ as reported in [167].

5.5 Data quality consideration

This section examines different data generation strategies and their impact on the performance of the DL model. The size and quality of the input dataset directly influence the predictive capability of the DNN, and performance is evaluated using MSE and R^2 metrics. The aim here is to underscore the critical role of both data quality and data quantity in shaping DL performance, thereby enabling informed decisions about the trade-offs between computational efficiency and predictive accuracy in practice. In Section 5.3.2, the network input was restricted to salient spectral features. While this representation significantly simplifies the network and reduces the dimensionality of the input space, it inevitably discards much of the rich information contained in the full spectrum. As a result, the model may overlook subtle yet important patterns such as bandwidth characteristics, or overall spectral shape, that are essential for accurately capturing the underlying geometry–response relationship. This simplification may subsequently trade accuracy for efficiency, and may have contributed to the sub-optimal predictive performance observed in the metrics presented in Table 5.1. In the subsequent investigation, 6 layers are used, with architecture: $\mathbf{u} \rightarrow 1024 \rightarrow 512 \rightarrow 256 \rightarrow 128 \rightarrow 64 \rightarrow 2$. Batch normalisation is applied after each layer to ensure stable and efficient learning. This architecture is used consistently across all three cases in the ensuing study to enable fair comparison. Early stopping with a patience of 50 epochs is implemented to prevent overfitting, and the ADAM optimisation algorithm is used for training. The outputs in all cases are p and Nt .

5.5.1 Peak information data

In this approach, the input vector to the DNN is a compact representation derived from the spectral data. Specifically, the peak absorption value, α_p , and its corresponding resonant

frequency, f_r , are extracted and used as input features, i.e., $\mathbf{u} = [\alpha_p, f_r, H]$. The network outputs are the absorber geometric parameters p , N_t , L_s , and H . By focusing on peak-based descriptors rather than the full 901D spectrum, the dimensionality of the input space is drastically reduced. This not only decreases the computational burden but also simplifies the learning task, as the network is trained to map a small set of salient acoustic features to the underlying design parameters. However, this feature selection also limits the amount of spectral information available to the model. Bandwidth characteristics, or subtle deviations across the frequency range are not explicitly captured, which may reduce predictive accuracy for more complex absorber behaviors.

5.5.2 Peak data enriched with bandwidth information

In this configuration, the input vector to the DNN is extended beyond the peak absorptivity and resonant frequency to include bandwidth information. Specifically, the upper and lower band-edge frequencies, f_u and f_l , are extracted from the spectral data using the full width at half maximum (FWHM) method. The resulting input vector is therefore defined as $\mathbf{u} = [\alpha_p, f_r, f_u, f_l, H]$. This representation increases the dimensionality of the input while embedding information on both the peak response and the effective absorption bandwidth. The inclusion of band-edge frequencies provides a richer description of the spectral response, capturing not only the strength and location of the resonance but also its spread across frequency. Such information is critical to distinguish spectral characteristics, which are not adequately represented when only the peak information is considered. Consequently, the model is expected to achieve improved accuracy in mapping from acoustic features to absorber geometry, particularly in cases where absorbers exhibit similar peak values but differ in their spectral widths. The same feedforward network configuration illustrated in Fig. 5.1 is employed.

5.5.3 Full spectrum data

In this configuration, the full 901D spectral response is used as input to the DNN. The use of high-dimensional spectral data preserves the complete frequency-dependent information of the absorber within the defined operational range, ensuring that subtle features across the absorption response are not lost through pre-processing or feature reduction. However, the large input dimensionality significantly increases the complexity of the learning problem. To capture the intricate, often nonlinear relationships between the spectral signatures and the

underlying geometric parameters, the network architecture must therefore be correspondingly deeper and more expressive, with sufficient capacity to avoid underfitting while mitigating the risk of overfitting.

Table 5.4 summarises the result of the comparison in terms of R^2 and MSE

Table 5.4: Performance of the DNN under different input feature configurations.

Input vector, $\mathbf{u}$	MSE	R^2
$[\alpha_p, f_r, L_s, H]$	3.89×10^{-3}	94.47%
$[\alpha_p, f_r, f_u, f_l, L_s, H]$	1.06×10^{-3}	98.48%
$[\alpha_1, \alpha_2, \dots, \alpha_{901}, L_s, H]$	1.85×10^{-3}	97.51%

Table 5.4 compares the predictive performance of the DNN across three different input scenarios. When only peak information (α_p, f_r) and size-related parameters (L_s, H) are used, the model achieves a moderate accuracy of $R^2 = 94.47\%$. Incorporating bandwidth descriptors (f_u, f_l) significantly improves performance, reducing the MSE to 1.06×10^{-3} and raising R^2 to 98.48%. Using the full 901D spectrum as input yields $R^2 = 97.51\%$, which is better than the peak-only case but slightly inferior to the bandwidth-enriched representation. This slightly lower performance may be attributed to redundancies and small variations in the full spectral data that may not directly relate to the geometry. This result highlights the trade-off between information richness and dimensionality: peak-only features are too limited, full-spectrum data is prone to overfitting despite its completeness, while bandwidth-enriched features strike the most effective balance, providing high accuracy with relatively low input dimensionality. Table 5.8 presents a visual assessment of the result. This outcome motivates the choice of bandwidth-enriched inputs for the autoencoder architecture used in the next chapter for robust AMM design synthesis.




Figure 5.8: Comparison of DNN performance under different input feature configurations. Three input scenarios are evaluated: (i) peak-based features, (ii) bandwidth-enriched features, and (iii) full-spectrum features. The bandwidth-enriched input achieves the best balance of accuracy and dimensionality, yielding lower MSE and higher R^2 compared to both peak-only and full-spectrum inputs.

5.6 Conclusions

In this chapter, a detailed performance comparison between deep learning (DNN) approaches and traditional optimisation (TO) methods for AMM design, was presented. Specifically, a DNN model with three hidden layers and one output layer was trained, and its performance was benchmarked against a traditional generalised pattern search (GPS) optimisation methodology. Multiple metrics, including robustness and time complexity, were employed to assess the relative performance of each approach. The study demonstrated that, once trained, a DNN offers a significant advantage in prediction efficiency, often producing design outputs in milliseconds, which is especially beneficial for complex AMM structures. However, the accuracy of DNN predictions can be comparable to that of TO algorithms, which inherently leverage the underlying physics of the system. Interestingly, the study also revealed that the predicted geometric or design parameters from DNN and TO may differ substantially, yet still yield the desired acoustic response, highlighting the existence of multiple viable design

solutions for the same target performance. While DNN models enable rapid exploration of non-intuitive designs, their generalisation to new operating conditions or design scenarios can be time-consuming. Re-training the model for each new scenario can be computationally intensive and may disrupt workflow. In contrast, TO methods, once the underlying physical model is established, can often be efficiently adapted to new design conditions, providing a more efficient approach to test/investigate new design cases.

Furthermore, the study investigated the impact of input feature configuration on DNN performance. It was observed that excessively redundant information can lead to overfitting and reduce model efficiency, whereas overly limited input may compromise predictive accuracy. A balanced, bandwidth-enriched feature set was found to provide the most effective compromise, achieving high accuracy with relatively low input dimensionality. This insight informs the design of the training data employed in subsequent chapters.

Overall, the main advantage of DNN lies in its computational speed and capacity to rapidly generate design solutions, while TO provides the advantage of quickly assessing the performance of new designs for on-the-fly application.

Chapter 6

Robust Design Considerations

6.1 Introduction

The use of approximate simulation models and the inherent imperfections introduced by additive manufacturing (AM) present persistent challenges to the reliable performance of acoustic metamaterials (AMMs). While these limitations are well recognised within the research community, systematic efforts to directly mitigate their impact remain limited. This chapter presents a deep learning framework to bridge the gap between what is designed, what is expected, and what is actually obtained. A Structured Latent Autoencoder (SLAE) is trained on a simulated dataset generated from the HRAM absorber model detailed in Chapter 3. To enhance robustness, two complementary integration strategies are proposed. The first involves a data augmentation scheme, wherein the original spectral dataset is perturbed using informed transformations, such as spectral scaling, shifting, noise injection, and bandwidth modulation. This strategy serves a dual purpose: it compensates for inaccuracies stemming from model simplifications or unmodeled physical effects, and it simulates spectral distortions induced by AM-related fabrication errors. The second strategy introduces a transfer learning pipeline that incorporates experimental data to adapt the pre-trained SLAE to real-world manufacturing deviations. This transfer learning approach enables the model to generalise beyond idealised simulations by aligning with empirical observations. Collectively, these methods establish a robust and data-efficient inverse design framework capable of producing AMMs that are not only high-performing in theory but also resilient to fabrication uncertainties in practice. By embedding robustness directly into the learning pipeline, this work advances the reliability and practical deployability of AI-assisted AMM design. This chapter is organised as follows. Section 6.2 outlines the motivation for interest in robust design methods in the context of AMM, and highlights the challenges introduced by additive manufacturing. Section 6.3 presents an analysis of the dataset obtained from the analytical model, providing insights into the statistical distribution and underlying structure of the data. Section 6.4 reports the results of the network training and data augmentation studies, accompanied by a detailed discussion of their implications. Section 6.5 introduces the transfer learning framework that trains the model on experimental data (which inherently captures the fabrication artefacts) for a robust network. This is followed by a robustness analysis that evaluates the performance of the transfer learning-enhanced model against networks trained without augmentation or transfer learning. Section 6.7 demonstrates the practical applicability of the robust network by exploring its behaviour under different input wave-

form configurations. Finally, Section 6.8 summarises the key findings and contributions of the chapter. The results presented in this chapter are based on the author’s paper submitted to the Journal of Sound and Vibration, titled: Modeling Acoustic Metamaterials with Manufacturing Uncertainties Using Deep Learning

6.2 Motivation for fabrication-robust design

AMMs promise unprecedented control over sound, but their performance can be highly sensitive to geometric precision and material fidelity. While additive manufacturing (AM) enables rapid prototyping of complex AMM geometries, it inevitably introduces imperfections, such as deviations in geometry, surface finish, or material uniformity, that can significantly shift or degrade acoustic responses [17, 18, 192, 195]. Because AMMs rely on sub-wavelength resonances, such discrepancies are magnified, often leading to measurable mismatches between predicted and fabricated absorbers [9, 161, 162]. While these challenges are well recognised, most AMM design studies focus on idealised models or demonstrate the advantages of deep learning (DL) over traditional optimisation without fully addressing robustness. As a result, the reliability of data-driven methods under realistic manufacturing variability remains uncertain. This gap poses a critical obstacle to the transition of AMMs from laboratory demonstrations to real-world noise-control applications. The motivation for this chapter is therefore to reposition fabrication variability as an intrinsic part of the design process. The SLAE model, described in Chapter 3, is trained for the subsequent analysis. The dataset, generated from the analytical model based on the procedure described in Section 3.6, is first analysed to visualise the distribution of the design space. This step is critical to understanding the underlying structure and variability of the input space, which could impact the accuracy of the model’s predictions.

6.3 Data visualisation

The distribution of peak absorptivity in Fig. 6.1(a) is heavily right-skewed, with a strong concentration of values near unity. This implies that a large portion of the theoretical samples exhibit high absorption, suggesting that the dataset may be biased toward near-perfect absorbers. Despite this, the lower absorptivity regime (< 0.2) is still represented, albeit sparsely, indicating that the model will be exposed to a wide but imbalanced range of absorptivity values.

Figure 6.1(b) displays the distribution of the peak frequencies. The majority of spectral samples have peak absorption within the 100–200 [Hz] range, with a pronounced concentration around 130–150 [Hz], indicating a strong design preference or resonant behavior within this band. The number of spectra peaking at higher frequencies decreases monotonically beyond 200 [Hz], suggesting fewer configurations at these ranges under the current design constraints. A notable feature is the sharp monotonic drop in sample density between 300–450 [Hz], suggesting a non-uniform coverage of the frequency space. Although the full 50–500 [Hz] range is populated, the concentration around specific bands implies a sampling bias that could influence the generalizability of learned models.

The rightmost panel in Fig. 6.1(c) shows the distribution of bandwidths, computed as the absolute difference between the right and left band edges. The bandwidths span a wide range from ≈ 2.5 [Hz] to 182 [Hz], with a mean of approximately 24.27 [Hz]. The distribution is strongly skewed toward narrower bandwidths, with a heavy concentration below 50 [Hz], and a long tail extending toward broader resonances. This suggests that narrowband spectral responses are significantly more common in the dataset, potentially reflecting the dominant physical modes in the parameter space.

The observed asymmetries and modal distributions in the dataset primarily result from the inherent geometric constraints imposed during the data generation process. These constraints lead to non-uniform coverage of the parameter space, with clustering around specific regions and skewed distributions. Such characteristics introduce a significant level of complexity and challenge for data-driven modeling approaches, necessitating models capable of capturing nonlinear relationships and accommodating the uneven representation of features within the data.




Figure 6.1: Histogram distributions of key spectral parameters in the training dataset. (a) Peak absorptivity showing a strong concentration in the high absorptivity region. (b) Peak frequency exhibiting noticeable non-uniformity and modal clustering. (c) Bandwidth distribution, revealing a strong skew toward narrowband responses.

6.4 Results and discussion

Figure 6.2 shows the learning curve (see Chapter 3 for the SLAE network and data acquisition procedure) of the evolution of the mean squared error (MSE) for both the training and validation sets over 395 epochs. The initial sharp decline in both losses indicates that the model rapidly learns the underlying structure in the data during the early training phase (approximately the first 100 epochs). After this initial drop, both curves gradually flatten out, with minor fluctuations in the validation loss likely due to sensitivities to validation batch statistics. The lack of divergence between the two curves indicates that the model has neither underfit nor overfit. The final MSE values for both training and validation are well below 2×10^{-3} , suggesting that the SLAE architecture achieves a good balance between model complexity and performance. Early stopping was triggered after 50 consecutive epochs without improvement, selecting the model at epoch 345 with a validation loss of approximately 6.0×10^{-4} . This demonstrates rapid convergence during training. Overall, training took approximately 2 hours.




Figure 6.2: Learning curves of the SLAE on training and validation datasets. Both training and validation losses converge closely to a minimum, indicating stable learning and no significant overfitting.

To further evaluate the generalization ability and performance of the network on the hold-out set, kernel density plots for the outputs, p and Nt , are shown in Figs. 6.3(a) and 6.3(b), respectively. The smooth, continuous curves represent the probability density of the errors for each output. Most of the error values (approximately 68%) fall within one standard deviation

of the mean (indicated by the blue dashed vertical lines), while the red lines represent the means of the distributions, which are sufficiently close to zero.

Additionally, the error distributions appear approximately symmetric around their respective means, suggesting low systematic bias and an overall Gaussian-like shape. However, the distribution for parameter p exhibits greater variability in the tail region. Furthermore, the regression plots in Figs. 6.3(c) and 6.3(d) for p and Nt , respectively, show that the predictions lie sufficiently close to the 45° reference line.




Figure 6.3: Kernel density plots of the errors for outputs p and Nt (top row), and linear regression plots comparing true versus predicted values (bottom row). The vertical dashed lines in the KDE plots indicate the mean and one standard deviation from the mean for each error data. The linear regression plots show the relationship between the true and predicted values for each output. Labels (a), (b), (c), and (d) correspond to each plot as described in the figure.

Table 6.1 summarizes the metrics used to evaluate the network’s performance for each output, while Fig. 6.4 shows sample predictions based on three randomly generated spectra from the analytical model. Overall, good agreements are observed in all cases. Note that the predicted spectra are recovered from their predicted spectral features using the Lorentz model in Eq. (3.26).

Table 6.1: Summary of the evaluation metrics used to assess the network’s performance.

	Encoder (Inverse model)	
	MSE	R^2 (%)
p	1.92×10^{-4}	99.57
N_t	8.59×10^{-4}	98.92
	Decoder (Forward model)	
	1.0×10^{-3}	98.49




Figure 6.4: Predicted absorption spectra from SLAE (black lines) compared to the true spectra (red dots) for 3 randomly generated samples from the analytical model. The plots illustrate the model’s ability to accurately reconstruct the spectral features.

To more effectively validate the performance of the network, random geometric instances are numerically simulated in COMSOL® Multiphysics 6.1 with parameters defined in Table 6.2

Table 6.2: Sample Geometric parameters of the simulated absorber with $t = 0.8$ [mm] and $t_p = 1.5$ [mm] in all cases

Sample no	p	N_t	L_s	H
1	5.5 [mm]	4	48 [mm]	16 [mm]
2	3 [mm]	2.5	25 [mm]	15 [mm]
3	2 [mm]	1.8	20 [mm]	10 [mm]

Figure 6.5 presents a comparison between the predicted and corresponding true spectra obtained from numerical simulations. Overall, the predicted SLAE curves closely match the corresponding ground truth. However, a noticeable amplitude and bandwidth mismatches




Figure 6.5: Predicted absorption spectra from SLAE (black lines) compared to the true spectra (red dots) for 3 random samples based on finite element simulations. The plots illustrate the model’s performance to previously unseen spectral characteristics.

are observed in Fig. 6.5(c). It is important to note that the SLAE was trained using theoretical data whose spectral characteristics differ slightly from those of the numerical model, particularly in terms of bandwidth [177]. Nonetheless, the model shows good generalization performance despite the fundamental differences between the data sources used for training and evaluation. One potential approach to improve on this domain gap is to apply data augmentation techniques that introduce controlled variability into the training set.

The learning curve of the SLAE based on the augmentation described in Chapter 3, is shown in Fig. 6.6. Only the training set was augmented, the validation and holdout sets consisted exclusively of real data. To ensure a fair comparison, the same network architecture and training settings used previously were retained. The learning curve over 256 epochs shows the training loss (black curve) with a smooth and consistent decrease, indicating steady optimization and effective learning on the training data. The validation loss (gray curve), although exhibiting some fluctuations, follows a general downward trend and stabilizes around a low value after approximately 40 epochs. Importantly, the validation loss remains consistently lower than the training loss throughout the training process. This behavior, while somewhat atypical, can be explained by the use of data augmentation during training. Augmented data introduces additional variability and noise into the training process, effectively regularizing the model and potentially making the training data more challenging than the un-augmented validation set. This often results in a validation loss that appears lower, especially when augmentation is applied aggressively or when the validation set is relatively clean. Despite the inherent noise introduced by augmentation, the model demonstrates strong generalization capabilities, as evidenced by the low and stable validation loss. The absence of a

divergence between training and validation curves further suggests that the model is neither underfitting nor overfitting. Early stopping was triggered after 207 epochs with a validation loss of 9.0×10^{-4} due to lack of improvements in learning after 50 consecutive epochs. The resulting metrics based on holdout set are presented in Table 6.3, validating the adequacy of performance.

Table 6.3: Summary of the evaluation metrics for the augmented dataset scenario. The results demonstrate good predictive performance of the model under this condition, with R^2 values exceeding 90% and sufficiently low MSE values

Encoder (Inverse model)		
	MSE	R^2 (%)
p	4.0×10^{-4}	99.11
Nt	1.37×10^{-3}	98.28
Decoder (Forward model)		
	1.31×10^{-3}	98.02




Figure 6.6: Learning curve of the SLAE trained with a mix of real and augmented data. The plot depicts the training and validation losses over successive epochs, highlighting the model's convergence behavior. The inclusion of augmented data introduces additional variability in the training and validation sets, which can enhance the model's generalization capabilities and mitigate overfitting. The relatively stable validation loss demonstrates the robustness of the model under these conditions.

The impact of augmentation on the SLAE network is investigated using the samples presented in Table 6.2, and the resulting spectra are compared with those from the unaugmented model, as shown in Fig. 6.5, and with numerical simulations. The augmented SLAE is henceforth referred to as SLAE-AUG. The results demonstrate significant improvement,

particularly in Fig. 6.5(c), where the mismatch is notably more pronounced without augmentation. It is noteworthy that the scaling and shifting operations are particularly beneficial in sparse regions of the input space, where the available data may not sufficiently capture local variations. These transformations introduce plausible variations that help the model generalize better in under-sampled areas. In contrast, densely populated regions already exhibit a wide range of variations, making additional augmentations less impactful. As a result, the marginal utility of these augmentation is significantly higher in sparse regions, where they effectively enrich the training data and help account for limitations or simplifications in the analytical model.




Figure 6.7: Comparison of spectra obtained from SLAE-AUG, the unaugmented SLAE, and numerical simulations. The results highlight the improvements in spectral accuracy with augmentation, particularly in regions where the mismatch between the unaugmented model and the numerical response is most pronounced.

6.5 Manufacturing uncertainty consideration

In this section, a transfer learning (TL) framework is developed to incorporate real-world manufacturing imperfections into the final design of the AMM. The objective is to leverage the inherent variability introduced during the 3D printing process to enhance the material’s robustness. Rather than predicting specific defects directly, the framework aims to model how such imperfections manifest in the final product and affect its acoustic performance.

The pre-trained SLAE-AUG model is fine-tuned using experimentally measured absorptivity curves. However, the size-related input features, L_s and H , are subject to manufacturing uncertainties. These uncertainties can arise from factors such as material shrinkage, mechanical tolerances, material inconsistencies, and environmental conditions. To account for these effects, L_s and H are perturbed within the known manufacturing uncertainty bounds,

as defined in Table 6.4, which correspond to the properties of the printing material and fabrication technology employed.

To enhance robustness during TL, this perturbation process is repeated $N = 10$ times per sample, generating plausible geometric variations around the original design specifications. Each experimental spectrum is thus paired with N perturbed versions of (L_s, H) , resulting in a more representative input distribution that captures the range of possible printed geometries. This increases the original 16 experimental samples to 160 samples. The perturbation logic is based on Eq. (6.1).

$$\begin{aligned} L_{s\text{perturbed}} &= L_s \cdot (1 - s) \cdot (1 + m_{L_s}) + \delta_{xy} \\ H_{\text{perturbed}} &= H \cdot (1 - s) \cdot (1 + m_H) + \delta_z, \end{aligned} \tag{6.1}$$

where L_s and H , are the original design specifications (nominal values), and

- $s \sim \mathcal{U}(0.002, 0.005)$: shrinkage factor
- $m_{L_s}, m_H \sim \mathcal{U}(-0.05, 0.05)$: material variation factor
- $\delta_{xy} \sim \mathcal{U}(-0.05, 0.05)$: XY tolerance
- $\delta_z \sim \mathcal{U}(-0.03, 0.03)$: Z tolerance (applied to H)

Table 6.4: Key uncertainty sources and their ranges for PLA and Prusa i3 MK3S+ machine

Uncertainty source	Range
Shrinkage	0.2% to 0.5%
Material inconsistencies	$\pm 5\%$
XY tolerance (0.4 [mm] nozzle)	± 0.05 [mm]
Z-axis tolerance (0.1 [mm] layer height)	± 0.03 [mm].

6.5.1 v-fold cross-validation strategy

To ensure robustness and generalisability of the fine-tuned model in the presence of limited experimental data, the TL process is conducted using a 5-fold cross-validation strategy. Given the small number of available experimental samples, this approach maximises data utilisation while providing an unbiased estimate of model performance. In each fold, the dataset is partitioned into 90% training and 10% validation subsets. The model is fine-tuned on the training set, and its performance is evaluated on the validation set. This process is repeated 5 times, with each subset serving once as the validation set. The final performance metrics

are then averaged across all folds to account for variance due to data partitioning. This methodology allows reliable assessment of the model’s predictive accuracy and its ability to generalise to unseen experimental samples, despite the limited data available.

An early stopping threshold of 50 epochs is employed to prevent unnecessarily long training time. The model is fine-tuned using a learning rate of 1×10^{-4} and a batch size of 8 samples. The weighting factors, ψ and β in Eq. (3.25) are equally set to 1. During the cross-validation process, the global best model, i.e., the one achieving the lowest validation loss across all folds, is selected and retained for evaluation. The average performance metrics across all folds, combining both training and validation, are as follows: $R^2 = 98.33\%$ and $\text{MSE} = 1.5 \times 10^{-3}$, averaged over the predictions of both p and Nt .

The SLAE with TL is hereafter referred to as SLAE-TL. An instance of mismatch between numerical, experiment and analytical from Chapter 4 is revisited here to assess the performance to model/manufacturing imperfections. These mismatches were attributed to the assumption of acoustically rigid spiral boundaries in the simulation, and unmodeled physical effects. The plot is brought forward and shown in Fig. 6.8(a) for validation purposes. In Fig. 6.8(b), the SLAE-TL response is overlaid on the same plot, showing excellent agreement with the experimental curve in both bandwidth and peak absorptivity. This result highlights the potential of TL at incorporating real-world manufacturing variations into the design process, enhancing model fidelity without requiring exhaustive experimental retraining. Beyond its academic significance, such capability is crucial for industry adoption, where predictive reliability and reduced prototyping cycles are paramount. By enabling accurate translation of numerical/analytical designs into experimentally realisable devices, the SLAE-TL framework offers a robust pathway toward the scalable deployment (arising from its ability to efficiently adapt pretrained models to a wide range of AMM designs and fabrication scenarios with minimal additional data and computation) of AMM technologies. The performance is expected to improve further and converge as additional experimental samples, capturing broader manufacturing variability and a wider range of artifacts across the target frequency spectrum, are incorporated into the training process. While the current SLAE architecture uses four fully connected layers in both the encoder and decoder, future work could systematically explore the impact of deeper architectures on TL performance. Increasing network depth could help ensure that the frozen layers effectively capture important spectral features, while fine-tuning the later layers without distorting previously learned representations. However,

this approach must be carefully balanced to avoid overfitting. Investigating this trade-off represents a promising direction to enhance the robustness and generalisability of the SLAE-TL framework.




Figure 6.8: Comparison of absorptivity spectra highlighting the impact of transfer learning in capturing manufacturing-induced uncertainties. (a) Analytical and numerical models show clear discrepancies with the experimental spectrum, particularly in peak absorptivity and bandwidth. (b) Incorporation of transfer learning via SLAE-TL substantially reduces this mismatch, yielding predictions that closely follow the experimental curve. This demonstrates the ability of the TL to address the gap between theoretical models and fabricated AMMs.

6.6 Robustness analysis

In this context, robustness is better defined as consistency of behavior and stability of performance under parametric variations, rather than exact agreement with theoretical predictions or laboratory results. Indeed, a system may achieve perfect agreement in simulation or under controlled experimental conditions, yet fail completely when exposed to real-world conditions laden with uncertainties. The essential criterion is that the key physical trends, such as resonance locations, absorption bandwidths, and overall response characteristics, remain reproducible and fall within acceptable error margins. This level of agreement is sufficient to validate the design methodology and to support its application in practical scenarios.

6.6.1 Robustness to parametric uncertainty

To systematically evaluate the robustness of absorber designs against manufacturing uncertainties, the predicted geometric parameters, p and Nt , from SLAE-TL and SLAE for the target response shown in Fig. 6.8, are perturbed according to a normal distribution defined by the uncertainty ranges in Table 6.4. For each perturbed geometry, the acoustic absorption is computed using the analytical model. From the ensemble of simulated spectra, the mean and standard deviation of absorption across the frequency range of interest are calculated, providing a direct visualization of the sensitivity of each design to manufacturing imperfections. Robustness is visualized using a $\pm 2\sigma$ prediction band around the mean, which encompasses approximately 95% of the individual perturbed samples. This approach offers a conservative estimate of variability, capturing most potential deviations due to manufacturing uncertainties.




Figure 6.9: Comparison of robustness between SLAE and SLAE-TL designs under FDM manufacturing perturbations. Shaded regions represent $\pm 2\sigma$ prediction bands across the frequency spectrum, and solid black lines indicate mean absorption.

As shown in Fig. 6.9, the prediction band from SLAE-TL is much narrower than that of SLAE, indicating more robust predictions, as expected.

6.7 Rapid on-demand design

In practical scenarios, access to precise absorption spectra, whether through full-wave simulations or experimental measurements, can be limited due to computational cost, fabrication constraints, or a lack of prior data. To address this limitation, the SLAE-AUG network is used in a rapid on-demand design setting, where the input is not a measured spectrum, but a user-specified target waveform representing the desired spectral response. These target inputs are abstract but practically meaningful approximations of real-world performance requirements. For example, Gaussian and triangular profiles can model broadband or narrowband absorber responses, while rectangular waveforms may represent idealized bandpass or band-stop filters. The goal is for the network to interpret these user-defined waveforms and generate physically realisable structural parameters that produce absorption spectra closely matching the input. Figure 6.10 presents the outcomes of this task, where the SLAE-AUG network is presented with four distinct target waveforms: rectangular, Gaussian, and triangular profiles. The resulting predicted absorption spectra align closely with the input profiles, underscoring the network’s ability to generalise and accurately infer inverse mappings even from abstract and idealised inputs.

These results highlight several key strengths of the SLAE-AUG model:

- **Flexibility:** The network is not restricted to a fixed class of absorption spectra but can dynamically adapt to arbitrary user-defined targets.
- **Robustness:** Despite the non-physical nature of the input waveforms, the network consistently produces realistic outputs.
- **Precision:** The generated responses demonstrate minimal deviation from the original input waveforms, indicating that the network’s latent representation retains the essential spectral features required for reconstruction.

This capacity for rapid prototyping empowers designers to specify spectral objectives directly, without requiring a complete design workflow involving trial-and-error or large-scale simulation. Ultimately, this positions the design framework as a powerful inverse design tool for robust AMM design.




Figure 6.10: Demonstration of the SLAE-AUG network’s performance in rapid, on-demand inverse design. The network is provided with user-defined input waveforms representing target spectral behaviors, including (a) rectangular, (b) Gaussian, (c) triangular, and (d) another rectangular profile. For each case, the predicted absorption spectra generated by the network closely follow the desired waveform, illustrating the model’s flexibility and accuracy even in the absence of experimentally or numerically derived input spectra.

6.8 Conclusions

This study presents a novel approach for addressing the inherent uncertainties in additively manufactured parts using deep learning techniques. The focus here is not to eliminate these uncertainties but to systematically account for them during the design process to ensure that the desired material behavior is achieved in real-world applications. A simple yet effective autoencoder-based framework, termed SLAE, was proposed, which takes critical spectral characteristics as input and inversely predicts the corresponding geometric parameters. For validation, the SLAE model was also used to reconstruct the original spectral features from

the latent geometric space, thereby demonstrating its capacity to both map and preserve essential information. Given the discrepancies observed between analytical predictions and experimental measurements, data augmentation was employed to compensate for unmodeled physical effects and approximations. The spectral features were perturbed within a carefully estimated range to reflect these mismatches, which significantly improved the model's performance when compared to the unaugmented SLAE. The augmented model, SLAE-AUG, was then further adapted to real-world manufacturing conditions using a transfer learning framework. Additive manufacturing processes introduce a variety of uncertainties, such as material inconsistencies, shrinkage, and machine tolerances, that can affect the final part geometry and, consequently, the material's acoustic performance. To account for these manufacturing artefacts, the 16 fabricated geometries were augmented within the bounds of known manufacturing uncertainties (based on PLA material and Prusa printing technology), resulting in a dataset expanded to 160 samples. This process effectively incorporated real-world defects into the training procedure, significantly enhancing the model's ability to predict material behavior under practical conditions. The results highlight the effectiveness of combining deep learning, data augmentation, and transfer learning in addressing manufacturing-induced uncertainties. The SLAE-TL model demonstrated significant improvements in robustness and predictive accuracy, effectively bridging the imperfections introduced by 3D printing and unmodeled physical effects in the absorber. The approach employed in this chapter introduces a methodology for designing reliable and robust AMM, particularly in the context of additive manufacturing.

Chapter 7

Interpretable Inverse Modelling for Geometry Discovery through Design Space Segmentation

7.1 Introduction

Acoustic metamaterials (AMMs) offer promising capabilities for noise control and wave manipulation, but their nonlinear and multidimensional design spaces often conceal clear relationships between geometry and acoustic performance. The previous chapter addressed the critical challenge of robustness, focusing on how real-world variability in manufacturing or material properties can affect AMM performance. While robustness-enhanced models improve stability and reliability, they do not inherently provide insight into why certain designs work or how specific geometric features influence acoustic behavior. As such, there remains a significant gap in understanding the structure of the design space itself, how performance emerges from geometry, and whether patterns or regimes exist that can guide engineering decisions.

To address this, the current chapter introduces a data-driven framework aimed at interpreting the AMM design space through unsupervised segmentation. By clustering a large dataset of simulated absorptivity spectra, generated via systematic variation of geometric parameters, distinct performance regimes were identified, and reveal how geometry influences acoustic response. This segmentation not only offers physical insight into the design-performance relationship, but also enables inverse mapping: predicting geometric configurations given a desired spectral target. Each cluster is characterised through representative spectral profiles and statistical summaries of associated geometry. A surrogate classifier is trained to predict cluster membership from a given spectral response, enabling reverse mapping from performance targets to spectral classes, and a nearest-neighbor heuristic (NNH) is used to retrieve design candidates from within the predicted subspace. Overall, this integrated framework, with inherent robustness property, offers a flexible and explainable alternative to black-box deep learning methods. It supports targeted design exploration and lays a foundation for embedding interpretability into machine learning-assisted metamaterial design workflows. This chapter is primarily based on the author’s published article [196] and the conference proceedings [197].

The structure of this chapter is organised as follows: Section 7.2 presents a brief overview of data-driven applications, and the rational for design interpretability. Section 7.3 introduces an unsupervised clustering methodology using k-means to segment the absorption response space. In Section 7.4, an ANOVA-based statistical analysis is conducted to identify the geometric design parameters most influential in driving cluster formation. Section 7.5 discusses

the design implications derived from the spectral signatures of the identified clusters, offering actionable insights for AMM optimisation. Section 7.6 presents the development of a surrogate classifier capable of predicting cluster membership from arbitrary input absorption spectra. A practical design case study, demonstrating the utility of the proposed segmentation framework for inverse design, is provided in Section 7.7. Section 7.8 presents a unified pipeline for the segmentation workflow, incorporating a heuristic scheme for automatic geometry retrieval based on the predicted cluster, along with illustrative design examples. A comparative table highlighting the differences and advantages of the proposed method over conventional deep learning approaches is also presented. A practical application of the methodology to room acoustics is demonstrated in Section 7.9. Finally, the conclusions are summarised in Section 7.10.

7.2 Data-driven applications in acoustics

The design of AMMs has advanced significantly over the past decade, driven by their extraordinary capabilities in wave manipulation, including negative refraction, sound absorption, and sub-wavelength control [110, 111]. These functionalities are tightly coupled with complex geometrical features, which pose challenges for both intuitive understanding and computational optimization. The design of AMMs remains a computationally expensive process. Their performance is governed by complex, often nonlinear relationships between geometric parameters and resulting spectral behaviors, posing significant challenges to designers seeking optimal configurations under physical and manufacturing constraints [3, 198]. Traditionally, the modeling and optimization of AMMs have relied heavily on physics-based simulations, such as the forward modeling finite element methods (FEM), to capture the intricate interplay between geometry and wave propagation phenomena [199]. While accurate, these approaches are computationally expensive and not interpretable, particularly when exploring large/complex parameter spaces or performing inverse design.

To address the limitations of traditional optimisation, data-driven approaches have increasingly been adopted for AMM design. Surrogate and inverse models based on supervised learning, convolutional networks, and autoencoder architectures have shown strong potential in predicting acoustic responses and mapping between geometry and performance. More recent advances with generative models (e.g., VAEs, GANs) and tandem networks further extend flexibility in sampling and inverse design, while data augmentation strategies have

been explored to account for uncertainties and enhance robustness. Despite these advances, many methods focus predominantly on predictive accuracy, often at the expense of interpretability, which is critical for trust, debugging, and design guidance [200]. Understanding why a particular design performs the way it does is often as important as knowing how well it performs, especially when navigating high-dimensional design spaces and dealing with imperfect manufacturing realities, or flexibility in exploring design alternatives. Interpretability is particularly crucial in physical systems like AMMs, where designers must reason about trade-offs and constraints, especially under manufacturing imperfections. Hybrid modeling approaches that combine data-driven techniques with physics-informed priors have been proposed [201], but interpretability remains an open challenge. Clustering and dimensionality reduction methods, such as t-distributed Stochastic Neighbor Embedding (t-SNE) and Uniform Manifold Approximation and Projection (UMAP), have been employed in other fields (e.g., materials science, genomics) to identify meaningful structure in high-dimensional spaces [202], but their use in AMM design is still nascent.

This study introduces a novel design-space segmentation procedure through clustering of simulated spectral data to reveal natural groupings of AMM behaviors. It aims to bridge the gap between data-driven modeling and physical interpretability, offering insights into how design parameters influence performance and how these insights can inform both optimisation and fabrication workflows. Here, an interpretable clustering-based framework is proposed to segment the AMM design space based on spectral similarity. Using k-means clustering applied to a dataset of absorption spectra derived from simulations of parametrically varied AMM geometries, distinct spectral behavior regimes are identified, and these identified clusters are related to underlying geometric features, and represent subspaces in the design space where geometries with similar spectral responses can be retrieved. The approach enhances explainability by integrating statistical analyses of design parameters with visualisations of cluster-specific spectral trends. This segmentation not only aids in understanding structure-performance relationships, but also supports more informed and efficient design optimisation, ultimately contributing to the development of robust, application-specific metamaterials.

7.3 Characterisation of the absorber system

In order to interpret the vast design space of the absorber, an unsupervised clustering technique is used to identify behavioural regimes in the frequency-absorptivity space. k-means

clustering is a robust and interpretable method for partitioning the data into distinct regions based on shared acoustic responses. It is a centroid-based clustering algorithm that minimises intra-cluster variance by assigning each data point to the nearest cluster centroid in Euclidean space [203]. Given the high-dimensional nature and size of the simulation dataset used in this thesis, k-means is particularly well-suited due to its computational efficiency and scalability. Moreover, its simplicity allows for intuitive post-clustering interpretation in the context of AMM geometry and absorptivity behavior.

From a design perspective, k-means enables the identification of latent structure within the simulation outputs, helping reveal recurring patterns such as resonant absorption zones, or transitional regions of mixed behavior. These clusters correspond to meaningful physical regimes and provide a reduced representation of the design space, aiding interpretability and geometry retrieval. To assess clustering performance, a silhouette score metric [204], which quantifies how well-defined the clusters are based on the compactness and separation of the cluster assignments, is used. This value ranges between -1 and 1, where a higher score indicates better-defined and more distinct clusters. A value close to 1 implies that the data point is well matched to its own cluster and well separated from neighboring clusters, while a score near 0 suggests overlapping clusters, and negative values indicate potential mis-classification.

To enable a more interpretable visual analysis of the high-dimensional spectral data, the original 901D absorption spectra are reduced to a compact 2D representation by extracting the peak absorptivity value and its corresponding resonant frequency for each sample. This reduction preserves key features of the spectra, while also facilitating meaningful clustering and visualization. Subsequently, k-means clustering is applied to this 2D feature space, where each point represents a spectrum. The optimal number of clusters, $k=5$, is selected based on grid search optimization using the silhouette score as the evaluation metric. This yielded a moderate silhouette score of approximately 0.47, indicating reasonably well-separated clusters with some overlap or transitional regions. The presence of gradual transitions and sparse sub-regions can naturally result in lower silhouette values [205]. The resulting segmentation, shown in Fig. 7.1, delineates distinct regions in the absorptivity–frequency plane, each corresponding to a characteristic response type. Each colour in the figure represents a distinct cluster region as identified in the 2D feature space. The clustering provides an intuitive map

of the design space and reveals how spectra with similar peak behaviors are grouped, offering insights into performance trends across the dataset.

The relatively low silhouette score may be attributed to smooth or non-distinct transitions within the dataset, likely stemming from the data acquisition process and the constraints imposed during sampling [190]. These effects cause certain data points to deviate significantly from coherent group structures, resulting in sparse or transitional regions that behave like noise from a clustering perspective. Such behavior degrades the compactness and separation of clusters, thereby lowering the overall silhouette score. Despite the moderate silhouette score, the clustering reveals meaningful segmentation, providing a basis for subsequent region-specific physical interpretation.




Figure 7.1: Cluster analysis of the absorber system using k-means clustering with $k = 5$. The data points represent peak absorptivity and the corresponding frequency extracted from the spectral response. Each cluster groups similar absorber configurations, revealing underlying trends in acoustic performance.

Based on the cluster labels assigned by the k-means algorithm, each sample in the dataset is associated with one of the five identified clusters. These labels are then used to map the original 901D spectral responses to their respective clusters. For each cluster, an average absorption spectrum is computed by taking the mean of all spectra belonging to that group.

This approach provides insight into the dominant acoustic behavior within each cluster, highlighting representative spectral features and variations across different design regimes. The resulting average spectra for each cluster are shown in Fig. 7.2. Figure 7.2(a) is the representative average absorption spectra for each of the five identified clusters. The curves reflect distinct resonant behaviors and bandwidth information, where each cluster captures a characteristic frequency-dependent absorptivity profile. This serves to evaluate whether clusters defined in a highly reduced feature space still reflect meaningful distinctions in the broader spectral domain. As shown in the figure, each cluster exhibits a distinct absorption profile, with variations in resonance frequency, peak sharpness, and bandwidth. For instance, Cluster 2 and Cluster 4, display peaks at low frequencies below 200 [Hz], while Cluster 0 is characterized by samples with mostly high absorptivities, and a dominant peak around 250 [Hz]. Samples in Cluster 1, which show the broadest absorption peaks, are more common at higher frequencies, whereas samples with the narrowest peaks dominate Cluster 2 at very low frequencies. These results suggest that even clustering based on peak information alone captures significant structural diversity among the spectra. The relatively low overall absorptivity levels, which is a natural consequence of averaging many individual spectra with varying peak locations and amplitudes, cause sharp resonance features to be smoothed out. However, distinct spectral behaviors are still observable. This further justifies the use of full-spectrum surrogate modeling for nuanced classification, as it incorporates the complete energy distribution and enables more accurate inference of geometric design classes. Figure 7.2(b) shows the distribution of samples across the clusters. Cluster 4 contains the largest proportion of samples. This suggests that a significant portion of the absorber designs are concentrated in this regime, potentially reflecting the physical design space or parameter constraints that favor low-frequency absorption, potentially biasing predictions toward this majority class. In contrast, the remaining clusters exhibit more balanced but comparatively smaller sample populations within the spectral space.




Figure 7.2: (a) Averaged absorptivity spectra corresponding to each of the five clusters, illustrating distinct spectral behaviors. (b) Histogram showing the number of samples per cluster, with Cluster 4 dominating the distribution, indicating a significant concentration of responses in the low-frequency range.

7.3.1 Design-space mapping of spectral clusters

To gain physical insight into the clustering outcomes and their implications for metamaterial design, it is essential to relate the identified spectral clusters back to the underlying geometric parameters of the absorber structures. While the clustering was performed in a reduced spectral space, specifically using the peak absorptivity and its corresponding frequency, the original dataset was generated from a physics-based simulation model in which each sample corresponds to a unique configuration of geometric design parameters. This section aims to reveal the structural organisation of the clustering outcome by projecting the cluster labels back onto the geometric design space. This analysis provides deeper insight into the relationship between design structure and acoustic response.

Each absorber configuration in the dataset is parameterised by a set of geometric descriptors, i.e., H , L_s , p , and N_t . By examining the distribution of these parameters across the assigned cluster labels, a reverse-mapping is performed from performance space to design space, thereby uncovering patterns that drive cluster formation. This approach provides a view of how regions in the design space give rise to localized behaviors in the spectral domain. More so, the projection enables the identification of dominant geometric features associated with specific acoustic regimes and provides actionable insights for guiding the inverse design

of AMMs prior to running computationally expensive finite element simulations or predictive modeling. By characterising the geometric composition of each spectral cluster, designers can leverage this knowledge to target desired absorption behaviors through informed parameter selection.




Figure 7.3: Boxplot distributions of geometric design parameters across the identified clusters. Each subplot corresponds to a specific parameter: (a) spacing between turns, p , (b) number of turns, N_t , (c) side length of cavity, L_s , and (d) absorber thickness, H . The boxplots reveal the intra-cluster variability and inter-cluster differences in parameter distributions, providing insights into how design parameters are grouped within each cluster. Notably, some parameters exhibit strong clustering trends, such as distinct separation in p and N_t , while others show overlapping ranges, indicating shared design characteristics.

Fig. 7.3 shows the distribution of design parameters across the five clusters. Each boxplot summarizes the spread and central tendency of each parameter within a cluster, providing insights into the structural regularities that define the clustering. In particular, Cluster 2 shows the lowest median and tightest spread in p , indicating consistently compact spiral

designs. Cluster 0 also has a relatively low median but with greater variability. Conversely, Clusters 1, 3, and 4 exhibit significantly higher median values, with Cluster 3 having the highest spiral spacing overall, implying more variability in the spiral spacing. Cluster 1 and Cluster 3 have relatively low median N_t values, implying fewer number of turns per design. These represent compact-channel designs better suited for high-frequency absorption, while Cluster 2 shows the highest median and widest spread, reflecting a concentration of designs with densely wound spirals and significantly varying. Designs here are more tuned toward low-frequency performance due to the longer spiral path that increases phase delay and viscous losses at lower frequencies. L_s exhibits a clear upward shift in median values for Clusters 2 through 4, which may indicate that wider cavities are more characteristic of the structures in Clusters 2–4. Clusters 0 and 1 contain designs with more compact cavity sizes, L_s , as reflected by their lower medians and interquartile ranges. H appears relatively consistent across all clusters, with overlapping interquartile ranges and similar medians. The lack of strong separation suggests that H may contribute less to the cluster differentiation relative to the other parameters.

In general, the inter-cluster variability in p and N_t is particularly significant, implying these parameters exert stronger influence in differentiating the acoustic behaviour encoded by the clusters. In contrast, the broader overlaps in H and L_s distributions suggest these features are less discriminative. Outliers in multiple clusters also highlight rare or edge-case geometries, potentially associated with performance anomalies or constraints in the design space, and may contribute to designs without a clearly defined cluster group. Overall, interpretable, physically meaningful patterns are revealed in the design parameterization, reinforcing the connection between structural configuration and emergent acoustic behavior.

7.4 Discriminative analysis based on ANOVA

To quantitatively assess the contribution of each geometric parameter to the formation of spectral clusters, a one-way analysis of variance (ANOVA) was performed on the clustered dataset. Unlike traditional ANOVA, which tests a factor across multiple groups simultaneously, a one-vs-rest strategy is considered here. For each cluster, the distribution of individual design parameters was statistically compared against the combined data from all other clusters, allowing the extent of each parameter’s discriminative power to be measured. The resulting F-statistics quantify how strongly each geometric feature contributes to inter-cluster

separation. These values are compiled and visualized in the heatmap of Fig. 7.4, where the horizontal axis lists the geometric parameters, and the vertical axis represents the identified clusters. Color intensity reflects the F-value magnitude, with warmer tones indicating greater discriminatory strength. From this visualization and the corresponding statistics, the following key insights can be drawn:

- **Cluster 0:** L_s emerges as the most dominant factor, exhibiting the highest F-value ($\approx 39,984$), indicating it is the principal geometric determinant here, and hence most effective for tuning the spectral performance in this group. H shows moderate influence, whereas p demonstrates comparatively low discriminative power for this cluster.
- **Cluster 1:** Both N_t and L_s serve as the primary discriminants, with F-values of approximately 35,546 and 30,792, respectively. This suggests a co-influence of these parameters in defining the cluster's acoustic behavior. p contributes marginally, while H plays a negligible role.
- **Cluster 2:** A strong dominance of p is evident, with an F-value nearing 47,925. N_t contributes secondarily, while the markedly lower F-value of L_s implies a limited role in the formation of this cluster's spectral profile.
- **Cluster 3:** This cluster is most strongly influenced by p , which shows the highest F-value across all clusters ($\approx 79,423$). Additional contributions from N_t and H suggest a multi-parametric influence, though p remains the dominant feature.
- **Cluster 4:** The discriminative influence is determined by N_t and L_s , with respective F-values of approximately 39,787 and 23,509. p plays a secondary role, while H , despite its comparatively lower F-value, remains statistically significant and may indicate subtle but consistent design sensitivity.

Thus, the dominant parameters driving the clustering are the spiral spacing, p , and number of turns, N_t , as evidenced by their consistently high ANOVA F-values across multiple clusters. These parameters exhibit the greatest discriminatory power in separating the geometric configurations of the AMM designs. This trend is also observed in the boxplots of Fig. 7.3, suggesting that these parameters serve as primary levers for controlling the acoustic behavior of the absorber.




Figure 7.4: Heatmap of ANOVA F-values showing the statistical significance of each geometric parameter across identified clusters. Higher F-values indicate greater variance between a specific cluster and the remaining dataset for a given parameter, reflecting its discriminative importance in the clustering structure. The color scale highlights the magnitude of these differences.

7.5 Design implications based on cluster characteristics

The clustering results and accompanying statistical analysis provide actionable insights into the selection of geometric configurations tailored to desired system behaviors. By examining the dominant parameters within each cluster, one can infer the most suitable cluster alignment based on specific performance goals or constraints.

7.5.1 Cluster-specific geometric profiles

Each cluster represents a statistically distinct grouping of geometric configurations, corresponding to characteristic absorption behaviors. The dominant parameters within each cluster, identified through ANOVA F-tests and visualized using heatmaps and boxplots, highlight which geometric features most strongly influence the observed spectral characteristics. These features serve as levers for tuning performance. Table 7.1 summarizes the dominant geometric drivers and their likely impact on resonance behavior.

This cluster-wise profile analysis offers a practical, data-driven approach to inverse design. Designers can begin by selecting a target acoustic behavior, then identify the clusters that exhibit matching parameter profiles. Within each cluster, performance can be effectively

Table 7.1: Cluster-wise summary of dominant geometric features and their acoustic implications

Cluster	Dominant features (based on ANOVA F values in Fig. 7.4)	Spectral characteristics (based on Fig. 7.1)	Influence on resonance behavior
0	L_s (dominant), H (moderate), p (lowest)	Broad absorption peaks, mostly high absorptivities	L_s strongly impacts resonance behavior
1	N_t , L_s (very strong), p (moderate)	Broadest absorption peaks	Resonance behavior strongly influenced by N_t and L_s
2	p (very strong), N_t , H (moderate)	Narrowest absorption peaks, mostly low absorptivities	p strongly impacts resonance; N_t , H introduce secondary control
3	p , N_t (highly dominant), H (moderate)	Intermediate peak widths	p and N_t exhibits stronger influence in resonance tuning
4	N_t , L_s (very strong)	Narrow peaks	Resonance strongly influenced by N_t and L_s

tuned by focusing on the most influential geometric parameters, as summarized in Table 7.1. This approach not only streamlines the prototyping phase but also enhances design efficiency, interpretability and control over the design-performance relationship.

7.6 Surrogate classifier

To facilitate straightforward mapping of spectral profiles to distinct geometric design categories, a surrogate classification model is introduced to systematically link absorption spectra to clustered classes. This approach transforms the unsupervised clustering outcome into a tractable classification task, allowing the identification of spectral characteristics responsible for cluster assignment. Among various candidate classifiers, a Random Forest (RF) model is selected due to its inherent interpretability, robustness to overfitting, and capacity to handle high-dimensional input data without extensive pre-processing. Unlike linear models or black-box neural networks, RFs offer a structured ensemble of decision trees, where each split provides insight into the most discriminative features for classification. Here, the classifier is trained using the full-resolution absorption spectra (901D inputs) as features, with the cluster labels, derived from k-means clustering, as target outputs. The dataset is first partitioned into training and testing subsets, with 80% of the samples allocated for model training and the remaining 20% reserved for evaluation. This split ensures that the model's performance can be validated on unseen data, thereby assessing its generalizability and robustness

in predicting cluster membership for new absorption profiles. This formulation not only reinforces the validity of the clustering structure, but also enables backward interpretation: given a desired spectral response, one can infer the most likely cluster and, by extension, a representative class of geometric designs. The 901D input spectra are transformed into a significantly lower-dimensional representation using a dimensionality reduction technique. While several methods such as UMAP, t-SNE, and principal component analysis (PCA), can be used for this purpose, PCA is chosen here due to its linear nature and the interpretability of its resulting components. Unlike nonlinear approaches (e.g., UMAP and t-SNE), which are effective for visualization but do not yield components directly relatable to the original spectral features, PCA enables a direct mapping between principal components and the original variables. The optimal number of components is determined via a grid search, selecting the smallest number of components that retain at least 95% of the cumulative explained variance, thus preserving most of the original data's variance while ensuring dimensionality reduction. Figure 7.5 shows the result of the optimization, indicating that 22 principal components are sufficient to capture 95% of the cumulative variance in the data. Consequently, the original 901D spectral dataset is reduced to a 22D representation, significantly lowering the data complexity while preserving most of its variance. Importantly, the reduced features remain interpretable, as each principal component corresponds to a linear combination of the original spectral features.




Figure 7.5: Optimal search for the number of principal components. The figure illustrates the evaluation process used to determine the ideal number of principal components to retain for dimensionality reduction. By analyzing the variance explained across different numbers of components, the optimal balance between model complexity and accuracy is determined.

Next, the RF classifier, as depicted in Fig. 7.6, is trained based on the PCA components to map spectral profiles to cluster labels. The RF is configured with 200 decision trees (DT).




Figure 7.6: Schematic illustration of the Random Forest (RF) classifier architecture. The model consists of an ensemble of $N = 200$ decision trees, each trained on a bootstrapped subset of the training data with randomly selected feature subsets at each split. The RF aggregates the predictions from all constituent trees through majority voting to assign a final class label. This ensemble strategy enhances generalization, reduces overfitting, and improves robustness to noise in high-dimensional input spectra.

To evaluate the effectiveness of the classifier, its performance is assessed using a confusion matrix, as shown in Fig. 7.7. A classification accuracy of approximately 99% is obtained on the test set which shows the adequacy of the model. The diagonal elements (top-left to bottom-right) represent correct classifications, and off-diagonal values represent misclassifications. Overall, the matrix is highly diagonal, indicating very high accuracy. The following can be observed from the figure:

- Cluster 0 has minor misclassifications into clusters 2, 3, and 4 (35, 37, and 63 respectively).
- Some bi-directional, meaningful cross-talk between clusters 3 and 4, which is not unexpected if these clusters share geometric or spectral similarities.
- Cluster 3 shows the highest number of misclassifications, though still a minor fraction. Misclassified mostly as Clusters 0, 1, 2 and 4 (likely due to overlaps in spectral response).
- Cluster 4 has small spillover into clusters 0, 1, and 3.

Despite the overall high classification accuracy, a degree of cross-talk is observed specifically between clusters 3 and 4 in the confusion matrix. This phenomenon, characterized by

mutual mis-classification, suggests a degree of geometric or spectral similarity between these clusters. Specifically, the acoustic absorption profiles corresponding to clusters 3 and 4 may exhibit overlapping features, causing the classifier to encounter ambiguity when delineating their boundaries. Such overlap may indicate a design continuum, where small geometric variations result in only marginal spectral shifts. Rather than indicating model failure, this cross-talk can provide meaningful insight into the sensitivity and proximity of certain design regimes, reinforcing the importance of considering transitional behaviors in geometric tuning.




Figure 7.7: Confusion matrix illustrating the classification performance of the surrogate model in predicting geometric cluster labels from absorption spectra. The diagonal entries indicate correctly predicted samples for each cluster, while off-diagonal elements represent misclassifications. High values along the diagonal suggest strong model accuracy, with minimal cross-talk between clusters. This performance supports the utility of the surrogate model in reliably mapping spectral responses to design-relevant geometric groupings.

7.7 Design scenarios

To demonstrate the practical utility of the proposed segmentation framework, we present case studies that show how a designer can leverage the analysis to derive suitable geometric configurations from a target absorption spectrum.

7.7.1 Design objective

Assume a designer aims to fabricate a structure that exhibits strong acoustic absorption within a narrow frequency band, specifically between 140–160 [Hz]. The target absorption spectrum, shown in red, in Fig. 7.9, is presented over the 130–180 [Hz] range. The trained classifier is used to predict the most probable geometric cluster that aligns with this spectral response. That said, the input spectrum is predicted as belonging to **Cluster 0**, a group characterized in Fig. 7.1 as having high absorptivities and mostly influenced by L_s . With the cluster label determined, the designer now refers to the geometric characteristics of Cluster 0, informed by: 1) the boxplot of Fig. 7.3, which reveals the typical ranges of the geometric parameters within the cluster, 2) the ANOVA/heatmap analyses of Fig. 7.4, which indicates the parameter L_s as the most significant parameter to control the resonance behavior. The designer can use this information to select parameter values within the interquartile range corresponding to Cluster 0, while L_s may be prioritized to fine-tune performance. For example, p , Nt , and H , may be set at the median values of the cluster, as indicated by the horizontal lines in the corresponding boxplots for cluster 0, while L_s may be adjusted more freely to achieve the target response.

However, we acknowledge that manually selecting suitable combinations of parameter values from these distributions can be challenging, especially when many parameters are involved and interact in complex, nonlinear ways. An alternative approach is to perform a bounded traditional optimization over this subspace. While this approach is feasible, it introduces additional computational overhead and may compromise the inherent geometric or manufacturing constraints present in the training dataset. To address this issue, we introduce a unified pipeline for design-space segmentation and present two alternative methods for automated geometry retrieval based on predicted cluster labels.

7.8 Geometry retrieval from cluster ID

Owing to the challenge that may arise from recovering the final geometry from the predicted subspace, two approaches are proposed to facilitate the retrieval process: a non-parametric nearest-neighbor heuristic (NNH) and cluster-wise regression models (CRMs).

- **Nearest neighbor heuristic (NNH):** This performs a non-parametric retrieval by identifying the closest spectrum within the predicted cluster. The corresponding geometric parameters of the nearest m training samples are returned as candidate designs.

This retrieval approach allows multiple candidate designs to be returned for a given target spectrum, providing the designer with flexibility to select the geometric profile best suited to their application. Moreover, by restricting the search to the predicted cluster, the method maintains semantic consistency with the learned segmentation of the design space. The NNH approach is computationally efficient, requires no additional model training, and guarantees physically realizable geometries as it draws from the existing dataset. However, it is limited to retrieving previously encountered designs and does not generalize beyond the training data.

- **Cluster-wise regression models (CRMs):** Here, a separate regression model (e.g., a random forest regressor) is trained within each cluster to learn the mapping from spectra to geometric parameters. Once a spectrum is classified into a cluster, the associated regressor is used to predict the geometric configuration. This method has the advantage of generating unseen or interpolated designs within the cluster’s design subspace, offering greater flexibility. Thus, each cluster is summarized by a representative regression model that captures the local structure-performance relationship. Unlike the NNH approach, CRM produces a single geometry per spectrum and may come with computational overhead.

The two approaches represent a trade-off between interpretability and generative capability. The NNH method is more interpretable and favors fast, conservative retrieval of existing designs, whereas CRM enables exploratory design synthesis by learning mappings within each clustered region of the design space. For simplicity and transparency, NNH is adopted in all case studies considered in this paper. This retrieval method also preserves any geometric/-manufacturing constraints imposed on the training dataset. The segmentation-classification-retrieval framework is streamlined into a unified pipeline, as illustrated in Fig. 7.8. The encircled region in the figure corresponds to the segmentation-classification stage, while post-processing is handled by the NNH, which retrieves candidate geometries from the predicted subspace (cluster). In this pipeline, clustering is performed on the PCA-reduced features, in contrast to the previous analysis where only peak features extracted from the original spectral data were used for clustering.

Note that clustering can also be performed directly on geometric parameters rather than on spectral responses, as done in this work. However, geometric proximity does not necessarily guarantee similar acoustic behavior. Such geometry-based clustering may be more




Figure 7.8: Modularized pipeline of the segmentation framework. The workflow consists of: (1) applying PCA to compress the spectral data; (2) clustering the reduced representation using k-means to define discrete regions in the design space; (3) training an RF surrogate to predict the cluster label of new spectral inputs; and (4) retrieving m geometry candidates using the nearest-neighbor heuristics within the predicted cluster.

appropriate for forward modeling tasks, where the objective is to predict the acoustic response of a given geometry. Importantly, the clustering here is not intended solely to maximize compactness, but to segment the design space in ways that enable interpretation, prediction, and retrieval. Consequently, standard metrics such as silhouette score, which optimize compactness and separation, may not fully capture suitability for these downstream tasks. In some cases, a clustering with a lower silhouette score may actually yield better results for retrieval task.

Clustering on PCA-reduced spectral features was briefly investigated and, although $k = 37$ clusters was found to be optimal in this case based on silhouette score of ≈ 0.3 , relying solely on this metric may not necessarily be adequate for the retrieval task, and the most appropriate value of k may ultimately depend on the quality of the geometries retrieved for a given input spectrum. Initial experiments suggest that the influence of k on retrieval performance is relatively minor. For instance, both $k = 5$ from previous clustering (Fig. 7.1) and $k = 37$ based on PCA-reduced features, yielded the exact same set of geometries. However, one may argue that the input spectrum used in this test lies within a well-defined (i.e., consistently clustered) region of the design space. As such, this observation may not yet be generalizable across the entire space. Furthermore, interpretability of the proposed segmentation framework is influenced by k . Too few clusters risk merging distinct spectral regimes, obscuring clear geometry–performance relationships, while too many clusters may fragment the design space into overly specific groups with limited generalizability. An intermediate k balances resolution and sample diversity, allowing dominant geometric trends to emerge within each

cluster. Consequently, the optimal k for interpretability is not necessarily the value that maximizes standard clustering metrics, but rather the one that best preserves meaningful structure for downstream prediction and geometry retrieval.

For simplicity and consistency with the case study in Section 7.7, the value of $k = 5$ is maintained. We then retrieved $m = 5$ candidate geometries corresponding to the target response shown in Fig. 7.9, with the resulting geometries (G) summarized in Table 7.2.

Table 7.2: Geometric parameter values of the 5 retrieved geometries corresponding to the target spectrum within the predicted Cluster 2.

	p	Nt	Ls	H
G1	2.11 [mm]	2.29	42.63 [mm]	14.79 [mm]
G2	2.32 [mm]	3.44	31.15 [mm]	16.29 [mm]
G3	2.6 [mm]	1.94	49.48 [mm]	18.91 [mm]
G4	2.68 [mm]	2.32	42.998 [mm]	18.87 [mm]
G5	2.77 [mm]	2.29	43.53 [mm]	19.75 [mm]

It is important to note that the scatter plot in Fig. 7.1 is based on only two summary features extracted from each spectrum: the maximum absorptivity and its corresponding frequency. Clustering in this reduced feature space simplifies the spectral representation by ignoring details beyond the peak, enabling intuitive visualization of the design space. However, in the unified framework shown in Fig. 7.8, clustering is performed on PCA-reduced features which is expected to improve classification performance due to enhanced representational fidelity.

While Euclidean distance in PCA-reduced space, d_{PCA} , provides a practical and efficient heuristic for spectrum retrieval, it may under-emphasize critical features such as peak amplitude or location, as it treats all points equally, i.e., every frequency bin contributes the same to the overall distance. To better capture more relevant features, a composite distance score that incorporates both global similarity (via PCA) and domain-specific peak characteristics, is introduced. The scoring function aggregates spectral similarity and physical feature alignment via a weighted sum. The resulting scores are sorted in ascending order, and the top m geometries corresponding to the lowest scores are returned as the best matches.

$$\text{score} = \gamma \cdot d_{PCA} + \beta \cdot |f_p^{\text{target}} - f_{p,i}| + \zeta \cdot |\alpha_p^{\text{target}} - \alpha_{p,i}|, \quad (7.1)$$

where f_p is the peak location, α_p is the peak amplitude at resonance, and the subscript 'target' and 'i', refer to the query spectrum and the i -th candidate spectrum in the predicted




Figure 7.9: The target absorption spectrum (shown in red circles) is provided as input to the NNH module depicted in Fig. 7.8. The classifier predicts the most appropriate cluster label from the segmented design space, which is then used by the NNH to retrieve candidate geometries. The curves labeled G1 through G5 represent the simulated absorption responses of the retrieved geometries listed in Table 7.2.

cluster, respectively. γ emphasizes spectra similarity in PCA-reduced space, β emphasizes on matching peak frequency, and ζ on matching peak amplitude. Setting all weights to 1 gives equal importance to global PCA-based similarity, frequency alignment, and peak amplitude matching.

In practical scenarios, simulated responses from analytical or numerical models are often not available as inputs. Therefore, alternative ways of representing the desired acoustic response must be used. These alternatives should capture the essential spectral characteristics of the target response to be effective. In this work, we investigate the use of synthetic input waveforms, specifically Gaussian, triangular, and rectangular shapes, as proxies for target spectra.

Figs 7.10(a), 7.10(b), 7.10(c), and 7.10(d) further demonstrate the pipeline’s flexibility by using different inputs types, highlighting its adaptability in handling diverse spectral shapes.

A key observation across these tests is the influence of the scoring parameters on retrieval performance. For all cases, weights were set as $\gamma = 1$, $\beta = 0.1$, and $\zeta = 3$, prioritizing peak amplitude alignment. Notably, setting $\gamma = 0$ in Figs 7.10(a) and 7.10(c) shifts the retrieval toward broader responses that better capture the broad shape of the input spectrum. The pipeline successfully processed these target spectra, despite not being drawn from the original training data, and returned geometries whose simulated responses closely matched the inputs. This illustrates the method's ability to generalize beyond training spectra characteristics, and adapt to diverse, user-defined spectral goals. Moreover, by tuning the scoring weights, designers can prioritize features such as peak amplitude or spectral shape, further enhancing design control.




Figure 7.10: Performance of the segmentation framework to varying input types. The predicted acoustic responses are based on analytically simulated geometric parameters retrieved from the NNH module. This allows to validate the accuracy of the retrieved parameters. All cases use scoring weights $\gamma = 1$, $\beta = 0.1$, and $\zeta = 3$, unless otherwise specified.

While the segmentation-based method is efficient and interpretable, its performance is sensitive to data coverage and clustering quality; insufficient sampling or ill-defined clusters

Table 7.3: Comparison between deep learning and segmentation-based methods for inverse modeling.

Criterion	Deep Learning (e.g., MLP, CNN, VAE)	Segmentation-based (PCA + Clustering + RF)
Training time	High – minutes to hours; often requires GPU	Low – seconds to a few minutes (≈ 3 [min] in this work); CPU is sufficient
Inference time	Fast – milliseconds	Fast – milliseconds to sub-seconds (≈ 0.7 [s] in this work)
Prediction accuracy	High – yields accurate result with enough data and tuning	High – retrieved geometries produce spectra closely matching the target when clusters are well formed
Interpretability	Low – black-box behavior	High – transparent, cluster-based pipeline
Model complexity	High – many layers, parameters, and hyperparameters; harder to debug and maintain	Low – modular structure with interpretable components; easier to maintain and adjust
Constraint preservation	Needs explicit constraint enforcement	Inherent – preserves training data structure
Generative capability	Yes – can produce novel, unseen designs	No – retrieval-based, not generative
Design flexibility	High – generalizes across design spaces	Moderate – constrained by training data distribution
Multi-solution retrieval	Limited – typically one output per input	High – returns multiple geometry options per input
Data efficiency	Low – needs large datasets	High – effective with fewer samples
Domain knowledge integration	Challenging – requires architectural embedding	Natural – reflected in clustering and scoring

may lead to poor retrieval accuracy. The following scenarios highlight key strengths of the developed framework:

- **Design interpretability:** Provides insight into why specific geometries are recommended.
- **Efficiency:** Avoids exhaustive forward simulations by directly mapping spectra to geometric clusters.
- **Flexibility:** Facilitates exploration of alternative designs within the geometric variability of a cluster.
- **Robustness:** Maintains similar spectral responses within a cluster, even under parametric variations.

7.9 Practical application to room acoustics

In this section, the segmentation framework is extended to the broadband case for room acoustic application. The aim here is to generate optimised broadband geometry that is

tailored to room-specific modal behavior and layout. The resulting design aims to maximise absorption efficiency within the 100–200 [Hz] range while minimising structural footprint and material usage. The generated geometry was evaluated through finite element simulations and validated against experimental reverberation time (RT60) measurements in a partially evacuated office environment. Performance is benchmarked against a commercial 10 [cm] thick Spektra A10 absorber. The way sound is perceived within a room results from both direct sound, traveling straight from the source to the listener, and reflected sound, which bounces off various surfaces. Key contributors to these reflections include architectural features such as walls, floors, and ceilings, but other objects, including furniture and even people, can also influence sound propagation. How sound interacts with a surface depends on factors such as the material composition and the angle at which the sound wave strikes it. The acoustical analysis of a room is essential for understanding and optimizing the behavior of sound within enclosed spaces for optimal sound experience. The interaction of sound waves with surfaces, objects, and air affects how sound is perceived, making room acoustics a critical factor in designing environments such as concert halls, theaters, recording studios, classrooms, offices, and residential spaces [206]. Poor acoustic design can lead to issues such as excessive reverberation, speech intelligibility problems, and unwanted echoes, which degrade the quality of auditory experiences. Artificial intelligence (AI) has emerged as a transformative tool in the field of acoustics, revolutionizing acoustic design by enabling advanced noise control solutions. With its ability to analyse vast amounts of data and predict sound behavior with high precision, AI allows for the creation of customized acoustic environments tailored to specific needs, empowering designers to achieve targeted sound control with greater efficiency, and achieving optimized acoustics in spaces ranging from open offices to concert halls, while reducing time and resources spent on traditional trial-and-error approaches. The two primary methods for modelling room acoustics are based either on solving the Helmholtz wave equation or its variants, or on the principles of geometrical acoustics which form the foundation of ray acoustics. Geometrical acoustics, a commonly used approach for room acoustic modeling [206–209], treats sound propagation as rays and disregards its wave nature. This assumption is accurate at high frequencies where the sound wavelength is much smaller than the objects and room features (e.g., walls, furniture). However, at low frequencies, wave phenomena such as diffraction and interference become more pronounced, leading to significant small-scale ef-

fects that geometrical acoustics cannot account for, thus necessitating wave-based methods such as the finite element modeling (FEM) or a hybrid approach [210].

7.9.1 Room description and numerical modelling


Here, the wave-based approach to modelling room acoustics is adopted using COMSOL multiphysics. The considered office space is a rectangular room of dimensions 580 [cm] $\times$ 266 [cm] $\times$ 280 [cm], with a window of dimensions 130 [cm] $\times$ 54 [cm] $\times$ 180 [cm] as depicted in Fig. 7.11(a). The office is constructed with solid concrete and brick surfaces with reflective acoustic characteristics. The walls are made of exposed bricks, offering high mass and rigidity. This results in strong sound reflection, with minimal sound absorption at lower frequencies. The surface texture of the brick may introduce some diffusion at mid-to-high frequencies. The ceiling consists of a solid reinforced concrete slab, providing a highly reflective surface that contributes to sound reverberation within the room. The lack of absorption may lead to longer decay times for higher frequencies unless treated with acoustic materials. The floor is a polished concrete surface, which is both acoustically reflective and structurally robust. The smooth finish enhances high-frequency reflections, potentially increasing room brightness. The window is a rectangular opening with a shallow cavity ending in the glass backing. The glass pane, set at the rear of the cavity, influences both sound reflection and absorption characteristics.

The acoustic response of the room is modeled using the lossless Helmholtz wave equation within the Pressure Acoustics Module of COMSOL[®] Multiphysics. The medium is excited with a monopole domain source, applied through a spherical volumetric source. This configuration approximates an omnidirectional point source, which is representative of a typical loudspeaker or sound source in enclosed spaces. The spherical region ensures uniform radiation of acoustic energy in all directions, minimizing directional bias and better simulating diffuse field conditions assumed in statistical models such as the Sabine and Eyring formulas.

Boundary conditions are assigned to the room surfaces and furniture based on their respective material properties, as detailed in Table 7.4, and their absorptivities across octave band center frequencies are obtained from the literature [211, 212]. The computational domain is discretized using a tetrahedral mesh, with a maximum element size of $\lambda/7$, where λ represents the wavelength corresponding to the highest frequency of interest.

Table 7.4: Absorption coefficients of the surface materials in the investigated office space.

Surface		Absorptivities in octave bands					
Type	Material	125 [Hz]	250 [Hz]	500 [Hz]	1000 [Hz]	2000 [Hz]	4000 [Hz]
Window	Glass	0.18	0.06	0.04	0.03	0.02	0.02
Window frame	Plywood	0.28	0.22	0.17	0.09	0.1	0.11
Ceiling/floor	Concrete	0.01	0.01	0.015	0.02	0.02	0.02
Door	Wood	0.1	0.07	0.05	0.04	0.04	0.04
Side walls	Brick	0.03	0.03	0.03	0.04	0.05	0.07
Table	Wood	0.15	0.11	0.10	0.07	0.06	0.07

**Figure 7.11:** Illustration of the office space drawing simulated in COMSOL. (a) shows the CAD drawing of the room with a table, book shelf and a drawer, and (b) shows the map of the speaker-microphone positions during measurements.

7.9.2 Reverberation time

Reverberation is a key acoustic parameter that defines how sound persists in a room after the source has stopped [213]. It significantly impacts both speech intelligibility and musical quality. Excessive reverberation can blur speech by masking short consonant sounds, while insufficient reverberation can make music sound lifeless and lacking in richness. The reverberation time, $RT60$, is the duration required for the sound level to decay by 60 dB after the source ceases. This metric provides insight into the room's acoustic properties and can be measured empirically by analyzing the sound pressure level over time. The $RT60$ in frequency domain is commonly computed according to the Sabine formula [214] in Eq. (7.2)

$$RT60 = \frac{0.161V}{A}, \quad (7.2)$$

where V is the volume of the room, and A is the total equivalent absorption area given by

$$A = \sum (S_i \cdot \alpha_i), \quad (7.3)$$

where S_i is the effective surface area of each material in the room, and α_i is the absorption coefficient of the material. The Sabine formula assumes that sound energy is uniformly distributed within the room, with minimal influence from objects and furniture on the overall acoustic response. However, in practice, perfect sound diffusion is rarely achieved, requiring adjustments in some cases to account for furniture or surfaces not accounted for due to the assumption.

7.9.3 Measurement setup

Reverberation time measurements are conducted in accordance with the ISO 3382-2:2008 standard. Two microphone positions, M_1 and M_2 are marked within the room, spaced at least 1 [m] from the nearest reflecting wall, and avoids symmetric positions. Three sound source positions are utilized: two located at opposite corners of the room (S_1 and S_3) and one at the center, S_2 (see schematic in Fig. 7.11(b)). The excitation signal is a time-domain logarithmic sine sweep, providing a high signal-to-noise ratio and enabling deconvolution to extract the room impulse response from which RT60 is estimated via Schroeder integration. The sweep is generated at -12 [dBFS] with a sampling rate of 44.1 [kHz].

A total of 3P_2 ($= 6$) source-microphone combinations were measured, and the results were averaged to obtain a single representative reverberation time curve. The office measurement setup comprises a Genelec 8020B active monitoring loudspeaker and an omnidirectional microphone, connected to the output and input channels, respectively, of a Focusrite Scarlett 2i4 audio interface. The measurement signals are generated and analyzed using Room EQ Wizard (REW) V5.31.3, allowing for precise estimation of the decay characteristics of the room impulse response.

7.9.4 Model validation and results

Figure 7.12 shows the one-third octave RT60 of the office space without any acoustic treatment. The measurements at each speaker-microphone position are shown in gray, while the average response is indicated in red (refer to Fig. 7.11(b) for speaker-microphone positioning). It is noteworthy that the office space was already in use, meaning some items could not be removed during measurements, which may have impacted the room's response. Subsequently, the numerical RT60 is compared to the experimental results. Initial simulation

shows that the numerical model tends to overestimate the reverberation time. This discrepancy is likely due to the presence of surfaces, crevices, cavities, and other elements within the office space that are not fully accounted for in the numerical model. To correct for this, the unknown absorption characteristics are estimated by calculating the absorption area required to match the difference between the numerical and experimental RT60 responses. This "deficit" in absorption area is used to generate frequency-dependent absorption data, which are subsequently applied to a section of the room's wall in the numerical model. The required absorption area is computed using Sabine's equation as

$$\Delta A(f) = A_{exp}(f) - A_{num}(f) = \frac{0.161.V}{RT_{60}^{exp}} - \frac{0.161.V}{RT_{60}^{num}}, \quad (7.4)$$

So that the correction absorption coefficient, $\Delta\alpha(f)$, is then obtained according to Eq. 7.5

$$\Delta\alpha(f) = \frac{\Delta A(f)}{A_{corr}}, \quad (7.5)$$

where A_{corr} is the correction absorption area. The result of this correction is shown in Fig. 7.13 which validates the numerical for subsequent analysis.




Figure 7.12: Acoustic response of the space measured using a logarithmic sine sweep excitation showing the reverberation time (RT60) and decay characteristics of the space.

Having validated the numerical room model against experimental measurements, a simulation study was conducted to evaluate the performance of a machine learning-based AMM absorber in comparison to a conventional industrial absorber, the Spektrum A10, which con-




Figure 7.13: Comparison between experimental and numerical reverberation times after correcting for the unknown absorption area in the space.

sists of a fabric cover and mineral wool core, with dimensions $60[\text{cm}] \times 60 [\text{cm}] \times 10 [\text{cm}]$. To achieve a desired broadband response, the segmentation framework (Fig. 7.8) was used to retrieve the geometric parameters for fabrication purposes. The target input spectrum was generated based on the room's reverberation response (Fig. 7.12) using a multi-Gaussian curve with randomised peak heights, and applied to the numerical model. For validation purposes, the predicted geometry from the NNH module was analytically simulated and the response compared to the desired input as shown in Fig. 7.14. Note that the absorption response corresponding to the predicted geometric parameters from the NNH module may also be recovered and used as the desired spectrum that is applied to the numerical model.

A fabricated sample of the AMM absorber is shown in Fig. 7.15(b). Each unit consists of a supercell comprising four inhomogeneous cells (described in Chapter 3), with each cell tuned to a different target frequency. In the numerical simulation, the corresponding absorptivities of the AMM are assigned to a $60[\text{cm}] \times 60 [\text{cm}]$ square area on the room surfaces and ceiling. The maximum thickness of the AMM units is 4 cm, in contrast to the Spektrum A10 (Fig. 7.15(a)), which has a thickness of 10 [cm].

For the private office space under consideration, it is desired to have RT60 between say $0.4 \leq \text{RT60} \leq 0.6 [\text{s}]$ at low frequencies ($< 200 [\text{Hz}]$) to improve speech intelligibility and reduce listener fatigue from reflected noise. Thus, a broad range of 100 to 200 [Hz] with absorptivities higher than 0.8 is presented to the encoder for which the corresponding




Figure 7.14: Comparison between the predicted response and the desired input curve, generated from 4 Gaussians with randomised peak heights, based on the region of interest as observed in the room’s reverberation response.




Figure 7.15: Feature size comparison of the absorber designs. a) a 60 [cm] × 60 [cm] Spektrum A10 industrail panel used as a benchmark, b) a sample fabricated structure of the AMM absorber

geometry is predicted, and the reconstructed absorption spectrum, $\alpha(f)$, is then applied to the absorber surfaces in the room model. For fair comparison, the AMM absorber is also designed with the same 60[cm] × 60 [cm] cross-section, but with a thickness of 3 [cm]. Figure 7.16 shows the performance comparisons of both absorbers. It can be observed that the AMM slightly improves the RT60 relative to the baseline in the range between 100 and 200 [Hz], but the effect at higher frequencies are very minimal. This is because the absorber is targeted for

the low frequencies (≤ 200 [Hz]) and more so, the range of influence of the absorber can be increased by increasing the number of inhomogeneous unit that makes up each broadband absorber, as low frequencies are characterized by sharp narrow peaks. The Spektrum A10 on the other hand performs better at higher frequencies (> 200 [Hz]) with lower RT60. Figure 7.17 shows a bar plot of the reduction in RT60 relative to the baseline. It is clear that the AMM provides a superior performance in the targeted low-frequency range, and beyond this range, the Spektrum A10 provides the better performance.




Figure 7.16: Performance comparison of RT60 curves from both absorbers relative to the baseline.




Figure 7.17: Relative difference in RT60 from the baseline from both absorbers in the targeted frequency range.

7.10 Conclusions

In this chapter, a framework for the design-space segmentation of AMMs, aiming to enhance interpretability and insight into the relationship between geometric parameters and spectral performance, was introduced. By clustering a large population of geometrically distinct AMM configurations based on their absorption spectra, latent structure can be uncovered in the design space that reflects fundamental modes of acoustic behavior. Through rigorous statistical analyses, including ANOVA and boxplot evaluations, the geometric parameters most responsible for driving the formation of each spectral cluster are identified. This enabled a principled understanding of how specific geometric features, such as spiral spacing, absorber thickness, or cavity width, influence performance across the frequency domain. To support practical design applications, a supervised surrogate model using a random forest classifier, was developed. This model accurately predicts a sample's spectral cluster from its frequency response alone, allowing reverse inference: from a desired absorption profile, one can now infer the most appropriate design region within the geometric space. Beyond the data-driven tools, the study introduced two designer-oriented components: (1) a design implications matrix, linking acoustic behavior categories to dominant geometric clusters, and (2) design considerations, providing actionable guidelines for translating desired performance characteristics into interpretable geometric modifications. Furthermore, to enable automatic geometry retrieval from predicted cluster labels, a non-parametric NNH approach was incorporated into a unified segmentation framework. This allows the retrieval of m user-defined candidate geometries corresponding to a given query (absorption spectrum) within the identified subspace. Furthermore, the effectiveness of the framework to different user-defined input types for real case-scenarios, was investigated.

The methodology was subsequently applied to design a data-driven acoustic AMM absorber aimed at mitigating low-frequency room acoustic issues, and compared its performance with a conventional industry-grade absorber. The targeted environment was an existing office space that was partially evacuated and acoustically characterised using the RT60 metric. To establish a baseline and validate the simulation framework, numerical modelling was carried out using the finite element method in COMSOL Multiphysics. Following validation, the segmentation pipeline was trained on a broadband spectral dataset and subsequently queried to predict geometries in the 100–250 [Hz] frequency range for the specified spectral response

for fabrication purposes. The desired spectral response was then applied to the numerical model to simulate the corresponding reverberation times.

The AMM absorber showed a promising performance within the targeted low-frequency range, with a modest reduction in RT60 compared to the untreated baseline response. However, the commercial Spektra A10 panel, with a thickness of 10 [cm], slightly outperforms the 3 [cm] thick AMM at higher frequencies above 200 [Hz], consistent with its broadband design for general acoustic use.

In summary, the methodology proposed in this chapter provided a transparent alternative to DL methods, and lays the foundation for intuitive and scalable AMM design, where insight is prioritised alongside performance. Rather than treating the design space as a black box, this segmentation-based approach makes the structure of the space visible, explainable, and ultimately more controllable, and the potential of integrating machine learning with acoustic design to develop lightweight, application-specific absorption solutions for modern architectural acoustics, was demonstrated.

Chapter 8

Conclusions

8.1 General conclusions

Additive manufacturing (AM), particularly 3D printing, has transformed the prototyping and fabrication of acoustic metamaterials (AMMs), enabling the realisation of intricate structures that were previously unmanufacturable. However, this advancement comes at a cost: the inherent imperfections and uncertainties introduced by AM processes, such as geometric deviations, surface roughness, and material inconsistencies, can severely compromise the acoustic performance and reliability of the resulting designs. To address this challenge, this thesis proposed data-driven frameworks that not only addressed fabrication-induced variability but also provided intuitive and physically meaningful explanations for predicted geometries. This interpretability component directly answers a key inverse design question: why a specific geometry is suggested in response to a given spectral input, enhancing both reliability and scientific insight of the design process. Two methodologies were proposed to this end: 1) a deep learning model that focused on addressing robustness using augmentation approaches and experimental sample measurement data, and 2) a machine learning framework that focused on providing insights/explanations into the inverse design process. The overarching objectives of this thesis were successfully realised through the following key contributions:

- **Development and validation of an AMM absorber (HRAE):** The HRAE absorber concept was developed, drawing inspiration from various designs reported in the literature, with emphasis on amenability to optimisation and manufacturing practicality, while retaining sufficient complexity for the ensuing data-driven studies. Its performance was validated through complementary numerical simulations, analytical modelling, and experimental measurements (Chapter 4).
- **Generation of a representative dataset and comparative analysis:** A comprehensive dataset was generated from the validated AMM absorber design space using a uniform sampling strategy, while ensuring adherence to geometric and manufacturing constraints. A comparative analysis was conducted to justify the preference for deep learning methodologies over traditional optimisation approaches, based on time complexity and robustness considerations (Chapter 5).
- **Development of a structured latent autoencoder (SLAE):** A simple but effective SLAE architecture was developed and trained to enable accurate forward and inverse

design predictions of AMM behaviour, thereby providing a powerful surrogate modelling framework (Chapter 6).

- **Incorporation of robustness to modelling and manufacturing uncertainties:** Novel strategies were proposed to mitigate the impact of manufacturing imperfections and modelling errors, including data augmentation techniques and a transfer learning framework that fine-tunes the pre-trained SLAE with experimental data. This substantially enhanced the robustness and generalisability of the data-driven approach (Chapter 6).
- **Introduction of an interpretable design-space segmentation framework:** A new design-space segmentation methodology was developed to provide interpretable and physically meaningful insights into the AMM design landscape. This approach offers a scalable, flexible, and computationally efficient alternative to conventional deep learning methods, while retaining scientific interpretability (Chapter 7).
- **Demonstration on a broadband room-acoustic absorber:** The segmentation-based framework was implemented to design a broadband AMM absorber tailored for room acoustic applications, demonstrating its effectiveness and practical relevance (Chapter 7).

8.2 Research contribution

This thesis makes a couple of impactful contributions to the field of AMM design, marking a significant step toward real-world deployability of data-driven metamaterial absorbers. It represents the first comprehensive effort to systematically integrate manufacturing-aware robustness, data-efficient deep learning strategies, and scientifically interpretable design tools within a unified AMM design framework. A major contribution of this work is the introduction of robustness to manufacturing uncertainties directly into the AMM design pipeline. While previous studies have often acknowledged discrepancies between analytical predictions, numerical simulations, and experimental measurements, attributing them to assumptions such as rigid-wall boundaries, geometric idealisations, and neglected loss mechanisms, they have largely treated such imperfections as unavoidable errors to be minimised post hoc. This thesis reframes these discrepancies not as obstacles but as informative signals that can be actively modelled and leveraged. By incorporating synthetic perturbations reflecting fabri-

cation tolerances into the training data, and by employing a transfer learning strategy to adapt pre-trained models using sparse experimental data, this work pioneers a data-driven pathway to proactively accommodate and compensate for model imperfections. This contribution establishes a principled route for bridging the persistent simulation–reality gap in AMM research. A second key contribution lies in the identification of optimal data representations for deep learning in AMM design. The thesis presents a comprehensive investigation of how different input configurations — from full absorption spectra to reduced peak-based descriptors — influence learning accuracy, and computational efficiency. This provides an actionable framework for selecting input formats based on task requirements, striking a balance between model complexity and predictive reliability. This insight addresses a critical yet previously overlooked challenge: how to best represent data for deep learning to achieve an optimal input configuration that balances model complexity and accuracy.

Furthermore, this work makes a paradigm-shifting contribution to interpretability in AMM design, an aspect largely missing from the current literature. By introducing a design-space segmentation approach, the thesis demonstrates that data-driven methods can yield scientifically meaningful insights, revealing the underlying structure–performance relationships within the AMM design space. This transforms deep learning from a black-box prediction tool into an interpretable design assistant, enabling designers to anticipate how geometric changes propagate through to acoustic performance. Such transparency is essential for establishing trust in machine learning-assisted design and accelerating industrial adoption.

Collectively, these contributions provide both the theoretical foundations and practical methodologies required to design robust, predictable, and optimised AMM absorbers ready for real-world implementation. The work offers the first cohesive framework that unifies physics-based modelling, robust deep learning, and interpretability-driven design exploration. It not only addresses fundamental gaps in the literature but also establishes a springboard for future research, where AMM design can evolve from isolated numerical exercises into a truly predictive, reliable, and industry-ready discipline.

8.3 Future work

While this thesis has made significant advances in the design, optimisation, and interpretability of AMMs, several promising avenues remain for further exploration. These directions build

upon the robust frameworks and methodologies developed herein and aim to push AMM research toward broader applicability, higher fidelity, and industrial relevance.

- **Multi-physics and multi-functional AMMs:** The current work focused primarily on airborne acoustic absorption. Future research could extend the frameworks to account for coupled phenomena such as structural vibrations, or fluid–structure interactions. Incorporating such multi-physics considerations would enable the design of multi-functional AMMs capable of simultaneously addressing sound attenuation, vibration isolation, and energy harvesting.
- **Physics-informed deep learning:** Integrating physical laws and governing equations directly into the learning models, such as through physics-informed neural networks or constrained autoencoders, could reduce data requirements and improve generalisation to untested regions of the design space. This approach would bridge the gap between empirical learning and first-principles theory, enhancing model reliability.
- **Formal uncertainty quantification and probabilistic robustness:** While this thesis addressed robustness via data augmentation and transfer learning, future work could develop formal uncertainty quantification frameworks. Techniques such as Bayesian neural networks or Monte Carlo propagation could systematically account for variability in manufacturing tolerances, material properties, and environmental conditions, leading to designs with provable reliability under uncertainty.
- **Experimental and full-scale validation:** While the present work validated designs at the sample level using impedance tube measurements, future studies should investigate full-scale implementations, such as room-scale acoustic absorbers or architectural integrations. This would provide practical verification and identify real-world performance factors not captured in small-scale experiments.
- **Interactive and real-time design tools:** Embedding the SLAE models, segmentation frameworks, and optimisation pipelines into interactive or digital twin platforms could enable real-time AMM design. Designers could explore “what-if” scenarios, assess trade-offs instantly, and iteratively refine designs with immediate feedback, dramatically accelerating development cycles.

- **Cross-domain generalisation:** The methodologies developed in this thesis could be extended beyond acoustic absorbers to other wave-based metamaterials, such as phononic, elastic, or vibroacoustic metamaterials. Such generalisation could significantly broaden the impact of this work, creating a unified data-driven framework for multi-domain metamaterial design.
- **Open-source tool development and community adoption:** Packaging the developed models, datasets, and optimisation frameworks into an open-source software environment would facilitate wider adoption, reproducibility, and collaborative advancement within the acoustics and metamaterials community.

Bibliography

- [1] Richard A Shelby, David R Smith, and Seldon Schultz. Experimental verification of a negative index of refraction. *Science*, 292(5514):77–79, 2001. doi: 10.1126/science.1058847.
- [2] Zhengyou Liu, Xixiang Zhang, Yiwei Mao, Yan Yang Zhu, Zhiyu Yang, Che Ting Chan, and Ping Sheng. Locally resonant sonic materials. *Science*, 289(5485):1734–1736, 2000. ISSN 0036-8075. doi: 10.1126/science.289.5485.1734.
- [3] Zhiyu Yang, Jun Mei, Min Yang, NH Chan, and Ping Sheng. Membrane-type acoustic metamaterial with negative dynamic mass. *Physical review letters*, 101(20):204301, 2008. doi: <https://doi.org/10.1103/PhysRevLett.101.204301>.
- [4] Yong Li, Xue Jiang, Rui-qi Li, Bin Liang, Xin-ye Zou, Lei-lei Yin, and Jian-chun Cheng. Experimental realization of full control of reflected waves with subwavelength acoustic metasurfaces. *Physical Review Applied*, 2(6):064002, 2014. doi: <https://doi.org/10.1103/PhysRevApplied.2.064002>.
- [5] Yifan Zhu, Xudong Fan, Bin Liang, Jianchun Cheng, and Yun Jing. Ultrathin acoustic metasurface-based schroeder diffuser. *Physical Review X*, 7(2):021034, 2017. doi: <https://doi.org/10.1103/PhysRevX.7.021034>.
- [6] Kun Tang, Chunyin Qiu, Manzhu Ke, Jiuyang Lu, Yangtao Ye, and Zhengyou Liu. Anomalous refraction of airborne sound through ultrathin metasurfaces. *Scientific reports*, 4(1):6517, 2014. doi: <https://doi.org/10.1038/srep06517>.
- [7] Yangbo Xie, Wenqi Wang, Huanyang Chen, Adam Konneker, Bogdan-Ioan Popa, and Steven A Cummer. Wavefront modulation and subwavelength diffractive acoustics with an acoustic metasurface. *Nature communications*, 5(1):5553, 2014. doi: <https://doi.org/10.1038/ncomms6553>.

- [8] Yong Li and Badreddine M Assouar. Acoustic metasurface-based perfect absorber with deep subwavelength thickness. *Applied Physics Letters*, 108(6), 2016. ISSN 1077-3118. doi: 10.1063/1.4941338.
- [9] Krupali Donda, Yifan Zhu, Shi-Wang Fan, Liyun Cao, Yong Li, and Badreddine Assouar. Extreme low-frequency ultrathin acoustic absorbing metasurface. *Applied Physics Letters*, 115(17), 2019. ISSN 1077-3118. doi: 10.1063/1.5122704.
- [10] Md Robiul Hossain, Eoghan P Ross, and Gareth J Bennett. Acoustic plate-valve resonator for low-frequency sound absorption. *AIP Advances*, 13(6), 2023. doi: <https://doi.org/10.1063/5.0142908>.
- [11] World Health Organization. Environmental noise guidelines for the european region. Technical report, World Health Organization. Regional Office for Europe, 2018. URL <https://www.euro.who.int/en/health-topics/environment-and-health/noise>.
- [12] Mathias Basner, Wolfgang Babisch, Adrian Davis, Mark Brink, Charlotte Clark, Sabine Janssen, and Stephen Stansfeld. Auditory and non-auditory effects of noise on health. *The Lancet*, 383(9925):1325–1332, 2014. doi: 10.1016/S0140-6736(13)61613-X.
- [13] Bridget M. Shield and Julie E. Dockrell. The effects of noise on children at school: A review. *Building Acoustics*, 10(2):97–116, 2003. doi: 10.1260/135101003768965960.
- [14] Jean-François Allard and Noureddine Atalla. *Propagation of Sound in Porous Media: Modelling Sound Absorbing Materials*. John Wiley & Sons, 2nd edition, 2009. ISBN 978-0-470-74772-3. doi: 10.1002/9780470747339.
- [15] Mohammad Javad Jafari, Mohammad Reza Monazam, and Marzieh Kazempour. Providing an optimal porous absorbent pattern to reduce mid to low-frequency sounds. *Journal of Environmental Health Science and Engineering*, 16:289–297, 2018. doi: 10.1007/s40201-018-0317-3.
- [16] Nicholas Fang, Di Xi, Jianhua Xu, Murat Ambati, Worasit Srituravanich, Cheng Sun, and Xiang Zhang. Ultrasonic metamaterials with negative modulus. *Nature Materials*, 5(6):452–456, 2006. doi: <https://doi.org/10.1038/nmat1644>.
- [17] John Kennedy, Lara Flanagan, Luke Dowling, GJ Bennett, Henry Rice, and Daniel Trimble. The influence of additive manufacturing processes on the performance of a

- periodic acoustic metamaterial. *International Journal of Polymer Science*, 2019(1): 7029143, 2019. doi: 10.1155/2019/7029143.
- [18] Agnieszka Ciochon, John Kennedy, Raphael Leiba, Lara Flanagan, and Mark Culleton. The impact of surface roughness on an additively manufactured acoustic material: An experimental and numerical investigation. *Journal of Sound and Vibration*, 546:117434, 2023. doi: 10.1016/j.jsv.2022.117434.
- [19] Hannah Hoppen, Hrishikesh Suresh, Felix Langfeldt, Wolfgang Gleine, and Otto von Estorff. Impact of surface roughness on acoustic resonators made of closed-cell lightweight foams. In *27th International Congress on Sound and Vibration, ICSV 2021*, 2021. doi: 10.1109/ICSV50964.2021.9563749.
- [20] Tomasz G Zieliński, Nicolas Dauchez, Thomas Boutin, Mikel Leturia, Alexandre Wilkinson, Fabien Chevillotte, François-Xavier Bécot, and Rodolfo Venegas. Taking advantage of a 3d printing imperfection in the development of sound-absorbing materials. *Applied Acoustics*, 197:108941, 2022. doi: 10.1016/j.apacoust.2022.108941.
- [21] Tomasz G Zieliński, Kamil C Opiela, Piotr Pawłowski, Nicolas Dauchez, Thomas Boutin, John Kennedy, Daniel Trimble, Henry Rice, Bart Van Damme, Gwenael Hannema, et al. Reproducibility of sound-absorbing periodic porous materials using additive manufacturing technologies: Round robin study. *Additive Manufacturing*, 36:101564, 2020. doi: 10.1016/j.addma.2020.101564.
- [22] Henning E von Gierke and Kenneth McK Eldred. Effects of noise on people. *Noise news international*, 1(2):67–89, 1993. doi: <https://doi.org/10.3397/1.3703002>.
- [23] George C Maling Jr. Noise. *Springer Handbook of Acoustics*, pages 1001–1057, 2014. doi: https://doi.org/10.1007/978-1-4939-0755-7_23.
- [24] Eulalia Peris et al. Environmental noise in europe: 2020. Technical report, European Environment Agency, 2020.
- [25] Stylianos Kephelopoulos, Marco Paviotti, Fabienne Anfosso-Lédée, Dirk Van Maercke, Simon Shilton, and Nigel Jones. Advances in the development of common noise assessment methods in europe: The cnossos-eu framework for strategic environmental noise mapping. *Science of the Total Environment*, 482:400–410, 2014. doi: <https://doi.org/10.1016/j.scitotenv.2014.02.031>.

- [26] Rainer Guski, Dirk Schreckenber, and Rudolf Schuemer. Who environmental noise guidelines for the european region: A systematic review on environmental noise and annoyance. *International journal of environmental research and public health*, 14(12): 1539, 2017. doi: <https://doi.org/10.3390/ijerph14121539>.
- [27] Erwan Nogues, Raphaël Le Bidan, and Dominique Pastor. Active noise control with digital pdm mems mics. In *2015 International Symposium on Consumer Electronics (ISCE)*, pages 1–2, 2015. doi: 10.1109/ISCE.2015.7177788.
- [28] Chuang Shi, Feiyu Du, and Qianyang Wu. A digital twin architecture for wireless networked adaptive active noise control. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:2768–2777, 2022. doi: 10.1109/TASLP.2022.3199992.
- [29] Cheng-Yuan Chang and Sheng-Ting Li. Active noise control in headsets by using a low-cost microcontroller. *IEEE Transactions on Industrial Electronics*, 58(5):1936–1942, 2011. doi: 10.1109/TIE.2010.2058071.
- [30] Dubravko Miljković. Brief introduction to active noise control in cars. In *2024 47th MIPRO ICT and Electronics Convention (MIPRO)*, pages 915–920, 2024. doi: 10.1109/MIPRO60963.2024.10569845.
- [31] Abdulkarim Shalool, Nasharuddin Zainal, Kok Beng Gan, and Cila Umat. An investigation of passive and active noise reduction using commercial and standard tdh-49 headphones. In *2016 International Conference on Advances in Electrical, Electronic and Systems Engineering (ICAEES)*, pages 606–609, 2016. doi: 10.1109/ICAEES.2016.7888118.
- [32] D. Cochrane, D.Y. Chen, and D. Boroyevic. Passive cancellation of common-mode noise in power electronic circuits. *IEEE Transactions on Power Electronics*, 18(3):756–763, 2003. doi: 10.1109/TPEL.2003.810858.
- [33] Martin Vasina, Katarina Monkova, Peter Pavol Monka, Drazan Kozak, and Jozef Tkac. Study of the sound absorption properties of 3d-printed open-porous abs material structures. *Polymers*, 12(5):1062, 2020. doi: <https://doi.org/10.3390/polym12051062>.
- [34] Lothar Cremer, Helmut A Müller, and Thomas D Northwood. Principles and applications of room acoustics, 1984.

- [35] P. Guignouard, M. Meisser, J. F. Allard, P. Rebillard, and C. Depollier. Prediction and measurement of the acoustical impedance and absorption coefficient at oblique incidence of porous layers with perforated facings. *Noise Control Engineering Journal*, 36:129–135, 1991. doi: <https://doi.org/10.3397/1.2827787>.
- [36] Malcolm J. Crocker. *Handbook of Noise and Vibration Control*. John Wiley & Sons, Inc., Hoboken, NJ, 2007. ISBN 978-0-471-39599-7. doi: 10.1002/9780470209707.
- [37] Dah-You Maa. Microperforated-panel wideband absorbers. *Noise control engineering journal*, 29(3):77–84, 1987. doi: <https://doi.org/10.3397/1.2827694>.
- [38] Dah-You Maa. Potential of microperforated panel absorber. *the Journal of the Acoustical Society of America*, 104(5):2861–2866, 1998. doi: <https://doi.org/10.1121/1.423870>.
- [39] Jian Kang and HV Fuchs. Predicting the absorption of open weave textiles and micro-perforated membranes backed by an air space. *Journal of sound and vibration*, 220(5): 905–920, 1999. doi: <https://doi.org/10.1006/jsvi.1998.1977>.
- [40] F-C Lee and W-H Chen. Acoustic transmission analysis of multi-layer absorbers. *Journal of Sound and vibration*, 248(4):621–634, 2001. doi: <https://doi.org/10.1006/jsvi.2001.3825>.
- [41] D.-Y. Maa. Theory and design of microperforated panel sound-absorbing constructions. *Scientia Sinica*, 18(1):55–71, 1975. doi: 10.1360/ya1975-18-1-55.
- [42] Dah-You Maa. Wide-band sound absorber based on microperforated panels. In *Proceedings of INTER-NOISE 1984*, pages 415–420. Institute of Noise Control Engineering, 1984.
- [43] Victor G Veselago. The electrodynamics of substances with simultaneously negative values of ϵ and μ . *Soviet Physics Uspekhi*, 10(4):509–514, 1968. doi: 10.1070/PU1968v010n04ABEH003699.
- [44] Rodger M Walser. Metamaterials: an introduction. *Introduction to complex mediums for optics and electromagnetics*, 1:295–316, 2003. doi: <https://doi.org/10.1117/3.504610.ch13>.

- [45] John B Pendry, AJ Holden, WJ Stewart, and I Youngs. Extremely low frequency plasmons in metallic mesostructures. *Physical review letters*, 76(25):4773, 1996. doi: <https://doi.org/10.1103/PhysRevLett.76.4773>.
- [46] John B Pendry, Anthony J Holden, David J Robbins, and William J Stewart. Magnetism from conductors and enhanced nonlinear phenomena. *IEEE transactions on microwave theory and techniques*, 47(11):2075–2084, 1999. doi: 10.1109/22.798002.
- [47] John Brian Pendry. Negative refraction makes a perfect lens. *Physical review letters*, 85(18):3966, 2000. doi: <https://doi.org/10.1103/PhysRevLett.85.3966>.
- [48] Costas M Soukoulis, Stefan Linden, and Martin Wegener. Negative refractive index at optical wavelengths. *Science*, 315(5808):47–49, 2007. doi: 10.1126/science.1136481.
- [49] Hao Ge, Min Yang, Chu Ma, Ming-Hui Lu, Yan-Feng Chen, Nicholas Fang, and Ping Sheng. Breaking the barriers: advances in acoustic functional materials. *National Science Review*, 5(2):159–182, 2018. doi: <https://doi.org/10.1093/nsr/nwx154>.
- [50] Sanjay Kumar and Heow Pueh Lee. The present and future role of acoustic metamaterials for architectural and urban noise mitigations. *Acoustics*, 1(3):590–607, 2019. doi: 10.3390/acoustics1030035.
- [51] Eli Yablonovitch. Inhibited spontaneous emission in solid-state physics and electronics. *Physical review letters*, 58(20):2059, 1987. doi: <https://doi.org/10.1103/PhysRevLett.58.2059>.
- [52] EY Yablonovitch. Photonic band-gap structures. *Journal of the optical society of America B*, 10(2):283–295, 1993. doi: <https://doi.org/10.1364/JOSAB.10.000283>.
- [53] Sajeev John. Strong localization of photons in certain disordered dielectric superlattices. *Physical review letters*, 58(23):2486, 1987. doi: <https://doi.org/10.1103/PhysRevLett.58.2486>.
- [54] Willem van Haeringen and Daan Lenstra. Analogies in optics and micro electronics. (*No Title*), 1990.
- [55] Sajeev John and Kurt Busch. Photonic bandgap formation and tunability in certain self-organizing systems. *Journal of Lightwave Technology*, 17(11):1931, 1999. doi: <https://doi.org/10.1103/PhysRevE.58.3896>.

- [56] Ertugrul Cubukcu, Koray Aydin, Ekmel Ozbay, Stavroula Foteinopoulou, and Costas M Soukoulis. Negative refraction by photonic crystals. *Nature*, 423(6940):604–605, 2003. doi: <https://doi.org/10.1038/423604b>.
- [57] Evan J Reed, Marin Soljačić, and John D Joannopoulos. Reversed doppler effect in photonic crystals. *Physical review letters*, 91(13):133901, 2003. doi: <https://doi.org/10.1103/PhysRevLett.91.133901>.
- [58] Chiyan Luo, Steven G Johnson, JD Joannopoulos, and JB Pendry. Subwavelength imaging in photonic crystals. *Physical Review B*, 68(4):045115, 2003. doi: <https://doi.org/10.1103/PhysRevB.68.045115>.
- [59] Michael M Sigalas. Elastic and acoustic wave band structure. *Journal of sound and vibration*, 158(2):377–382, 1992. doi: [https://doi.org/10.1016/0022-460X\(92\)90059-7](https://doi.org/10.1016/0022-460X(92)90059-7).
- [60] Michael Sigalas and Eleftherios N Economou. Band structure of elastic waves in two dimensional systems. *Solid state communications*, 86(3):141–143, 1993. doi: [https://doi.org/10.1016/0038-1098\(93\)90888-T](https://doi.org/10.1016/0038-1098(93)90888-T).
- [61] Manvir S Kushwaha, Peter Halevi, Leonard Dobrzynski, and Bahram Djafari-Rouhani. Acoustic band structure of periodic elastic composites. *Physical review letters*, 71(13):2022, 1993. doi: <https://doi.org/10.1103/PhysRevLett.71.2022>.
- [62] Zhengyou Liu, Che Ting Chan, and Ping Sheng. Analytic model of phononic crystals with local resonances. *Physical Review B—Condensed Matter and Materials Physics*, 71(1):014103, 2005. doi: <https://doi.org/10.1103/PhysRevB.71.014103>.
- [63] FR Montero De Espinosa, E Jimenez, and M Torres. Ultrasonic band gap in a periodic two-dimensional composite. *Physical Review Letters*, 80(6):1208, 1998. doi: <https://doi.org/10.1103/PhysRevLett.80.1208>.
- [64] Juan V Sánchez-Pérez, D Caballero, Rosa Martínez-Sala, Constanza Rubio, José Sánchez-Dehesa, Francisco Meseguer, Jaime Llinares, and F Gálvez. Sound attenuation by a two-dimensional array of rigid cylinders. *Physical Review Letters*, 80(24):5325, 1998. doi: <https://doi.org/10.1103/PhysRevLett.80.5325>.

- [65] WM Robertson and JF Rudy III. Measurement of acoustic stop bands in two-dimensional periodic scattering arrays. *The Journal of the Acoustical Society of America*, 104(2):694–699, 1998. doi: <https://doi.org/10.1121/1.423344>.
- [66] Valentin Leroy, Alice Bretagne, Mathias Fink, Herve Willaime, Patrick Tabeling, and Arnaud Tourin. Design and characterization of bubble phononic crystals. *Applied Physics Letters*, 95(17), 2009. doi: <https://doi.org/10.1063/1.3254243>.
- [67] R Lucklum and J Li. Phononic crystals for liquid sensor applications. *Measurement Science and Technology*, 20(12):124014, 2009. doi: 10.1088/0957-0233/20/12/124014.
- [68] MF Su, RH Olsson, ZC Leseman, and I El-Kady. Realization of a phononic crystal operating at gigahertz frequencies. *Applied Physics Letters*, 96(5), 2010. doi: <https://doi.org/10.1063/1.3280376>.
- [69] Ioannis E Psarobas. Phononic crystals—sonic band-gap materials, 2005.
- [70] F Bloch. Elektron im periodischen potential, bändermodell des festkörpers, z. *Phys*, 52:555, 1928. doi: 10.1007/BF01339455.
- [71] Charles Kittel and Paul McEuen. *Introduction to solid state physics*. John Wiley & Sons, 2018.
- [72] Pierre A. Deymier. *Acoustic Metamaterials and Phononic Crystals*, volume 173. Springer Science & Business Media, 2013. ISBN 978-3-642-31232-8. doi: 10.1007/978-3-642-31232-8.
- [73] Abdelkrim Khelif and Ali Adibi. *Phononic Crystals*. Springer, Berlin, Germany, 2015. ISBN 978-1-4614-9393-8. doi: 10.1007/978-1-4614-9393-8.
- [74] Abdelkrim Khelif, Abdelkrim Choujaa, Sarah Benchabane, Bahram Djafari-Rouhani, and Vincent Laude. Guiding and bending of acoustic waves in highly confined phononic crystal waveguides. *Applied physics letters*, 84(22):4400–4402, 2004. doi: <https://doi.org/10.1063/1.1757642>.
- [75] Toyokatsu Miyashita. Sonic crystals and sonic wave-guides. *Measurement Science and Technology*, 16(5):R47, 2005. doi: 10.1088/0957-0233/16/5/R01.

- [76] Jerome O Vasseur, A-C Hladky-Hennion, Bahram Djafari-Rouhani, F Duval, Bertrand Dubus, Yan Pennec, and PA Deymier. Waveguiding in two-dimensional piezoelectric phononic crystal plates. *Journal of applied physics*, 101(11), 2007. doi: <https://doi.org/10.1063/1.2740352>.
- [77] Abdelkrim Khelif, Abdelkrim Choujaa, Bahram Djafari-Rouhani, Mikael Wilm, Sylvain Ballandras, and Vincent Laude. Trapping and guiding of acoustic waves by defect modes in a full-band-gap ultrasonic crystal. *physical Review B*, 68(21):214301, 2003. doi: <https://doi.org/10.1103/PhysRevB.68.214301>.
- [78] F Van Der Biest, A Sukhovich, A Tourin, JH Page, BA Van Tiggelen, Z Liu, and M Fink. Resonant tunneling of acoustic waves through a double barrier consisting of two phononic crystals. *Europhysics Letters*, 71(1):63, 2005. doi: [10.1209/epl/i2005-10064-8](https://doi.org/10.1209/epl/i2005-10064-8).
- [79] A Khelif, Pierre A Deymier, Bahram Djafari-Rouhani, Jerome O Vasseur, and Leonard Dobrzynski. Two-dimensional phononic crystal with tunable narrow pass band: Application to a waveguide with selective frequency. *Journal of applied physics*, 94(3):1308–1311, 2003. doi: <https://doi.org/10.1063/1.1557776>.
- [80] Xiangdong Zhang and Zhengyou Liu. Negative refraction of acoustic waves in two-dimensional phononic crystals. *Applied Physics Letters*, 85(2):341–343, 2004. doi: <https://doi.org/10.1063/1.1772854>.
- [81] Manzhu Ke, Zhengyou Liu, Chunyin Qiu, Wengang Wang, Jing Shi, Weijia Wen, and Ping Sheng. Negative-refraction imaging with two-dimensional phononic crystals. *Physical Review B—Condensed Matter and Materials Physics*, 72(6):064306, 2005. doi: <https://doi.org/10.1103/PhysRevB.72.064306>.
- [82] Suxia Yang, John H Page, Zhengyou Liu, Michael L Cowan, Che Ting Chan, and Ping Sheng. Focusing of sound in a 3d phononic crystal. *Physical review letters*, 93(2):024301, 2004. doi: <https://doi.org/10.1103/PhysRevLett.93.024301>.
- [83] Hangyuan Lv, Xiaoyong Tian, Michael Yu Wang, and Dichen Li. Vibration energy harvesting using a phononic crystal with point defect states. *Applied Physics Letters*, 102(3), 2013. doi: <https://doi.org/10.1063/1.4788810>.

- [84] S Tol, FL Degertekin, and A Erturk. 3d-printed phononic crystal lens for elastic wave focusing and energy harvesting. *Additive Manufacturing*, 29:100780, 2019. doi: <https://doi.org/10.1016/j.addma.2019.100780>.
- [85] Guobiao Hu, Lihua Tang, Junrui Liang, Chunbo Lan, and Raj Das. Acoustic-elastic metamaterials and phononic crystals for energy harvesting: A review. *Smart Materials and Structures*, 30(8):085025, 2021. doi: 10.1088/1361-665X/ac0cbc.
- [86] Samira Amoudache, Rayisa Moiseyenko, Yan Pennec, Bahram Djafari Rouhani, Antoine Khater, Ralf Lucklum, and Rachid Tigrine. Optical and acoustic sensing using fano-like resonances in dual phononic and photonic crystal plate. *Journal of Applied Physics*, 119(11), 2016. doi: <https://doi.org/10.1063/1.4944600>.
- [87] Tianxi Jiang, Qingbo He, and Zhike Peng. Enhanced directional acoustic sensing with phononic crystal cavity resonance. *Applied Physics Letters*, 112(26):261902, 2018. ISSN 0003-6951. doi: 10.1063/1.5024970.
- [88] Oluwaseyi Ogun and John Kennedy. Noise attenuation through optimised acoustic metamaterials: A low form factor design for targeted noise reduction. In *2022 28th International Workshop on Thermal Investigations of ICs and Systems (THERMINIC)*, pages 1–4. IEEE, 2022. doi: 10.1109/THERMINIC57263.2022.9950667.
- [89] Jensen Li and Che Ting Chan. Double-negative acoustic metamaterial. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 70(5):055602, 2004. ISSN 1550-2376. doi: 10.1103/PhysRevE.70.055602.
- [90] Badreddine Assouar, Bin Liang, Ying Wu, Yong Li, Jian-Chun Cheng, and Yun Jing. Acoustic metasurfaces. *Nature Reviews Materials*, 3(12):460–472, 2018. doi: <https://doi.org/10.1038/s41578-018-0061-4>.
- [91] Nanfang Yu, Patrice Genevet, Mikhail A Kats, Francesco Aieta, Jean-Philippe Tetienne, Federico Capasso, and Zeno Gaburro. Light propagation with phase discontinuities: generalized laws of reflection and refraction. *science*, 334(6054):333–337, 2011. doi: 10.1126/science.1210713.
- [92] Yadong Xu, Yangyang Fu, and Huanyang Chen. Planar gradient metamaterials. *Nature Reviews Materials*, 1(12):1–14, 2016. doi: <https://doi.org/10.1038/natrevmats.2016.67>.

- [93] Yong Li, Bin Liang, Zhong-ming Gu, Xin-ye Zou, and Jian-chun Cheng. Reflected wavefront manipulation based on ultrathin planar acoustic metasurfaces. *Scientific reports*, 3(1):2546, 2013. doi: <https://doi.org/10.1038/srep02546>.
- [94] Jun Mei and Ying Wu. Controllable transmission and total reflection through an impedance-matched acoustic metasurface. *New Journal of Physics*, 16(12):123007, 2014. doi: [10.1088/1367-2630/16/12/123007](https://doi.org/10.1088/1367-2630/16/12/123007).
- [95] Guancong Ma, Min Yang, Songwen Xiao, Zhiyu Yang, and Ping Sheng. Acoustic metasurface with hybrid resonances. *Nature materials*, 13(9):873–878, 2014. doi: <https://doi.org/10.1038/nmat3994>.
- [96] Shuibao Qi, Yong Li, and Badreddine Assouar. Acoustic focusing and energy confinement based on multilateral metasurfaces. *Physical Review Applied*, 7(5):054006, 2017. doi: <https://doi.org/10.1103/PhysRevApplied.7.054006>.
- [97] Y Cheng, C Zhou, BG Yuan, DJ Wu, Q Wei, and XJ Liu. Ultra-sparse metasurface for high reflection of low-frequency sound based on artificial mie resonances. *Nature materials*, 14(10):1013–1019, 2015. doi: <https://doi.org/10.1038/nmat4393>.
- [98] Boyang Xie, Kun Tang, Hua Cheng, Zhengyou Liu, Shuqi Chen, and Jianguo Tian. Coding acoustic metasurfaces. *Advanced Materials*, 29(6):1603507, 2017. doi: [10.1002/adma.201603507](https://doi.org/10.1002/adma.201603507).
- [99] Shu-Yu Zuo, Ye Tian, Qi Wei, Ying Cheng, and Xiao-Jun Liu. Acoustic analog computing based on a reflective metasurface with decoupled modulation of phase and amplitude. *Journal of Applied Physics*, 123(9), 2018. doi: <https://doi.org/10.1063/1.5004617>.
- [100] Marc Dubois, Chengzhi Shi, Yuan Wang, and Xiang Zhang. A thin and conformal metasurface for illusion acoustics of rapidly changing profiles. *Applied Physics Letters*, 110(15), 2017. doi: <https://doi.org/10.1063/1.4979978>.
- [101] Jun Mei, Xiujuan Zhang, and Ying Wu. Ultrathin metasurface with high absorptance for waterborne sound. *Journal of Applied Physics*, 123(9), 2018. doi: <https://doi.org/10.1063/1.5009382>.

- [102] Yabin Jin, Bernard Bonello, Rayisa P Moiseyenko, Yan Pennec, Olga Boyko, and Bahram Djafari-Rouhani. Pillar-type acoustic metasurface. *Physical Review B*, 96(10):104311, 2017. doi: <https://doi.org/10.1103/PhysRevB.96.104311>.
- [103] Jun Lan, Yifeng Li, Yue Xu, and Xiaozhou Liu. Manipulation of acoustic wavefront by gradient metasurface based on helmholtz resonators. *Scientific reports*, 7(1):10587, 2017. doi: <https://doi.org/10.1038/s41598-017-10781-5>.
- [104] Jie Zhou, Xin Zhang, and Yi Fang. Analytical modelling for predicting the sound field of planar acoustic metasurface. *Journal of Applied Physics*, 123(3), 2018. doi: <https://doi.org/10.1063/1.5000055>.
- [105] Tuo Liu, Shanjun Liang, Fei Chen, and Jie Zhu. Inherent losses induced absorptive acoustic rainbow trapping with a gradient metasurface. *Journal of Applied Physics*, 123(9), 2018. doi: <https://doi.org/10.1063/1.4997631>.
- [106] Xue Jiang, Bin Liang, Xin-ye Zou, Jing Yang, Lei-lei Yin, Jun Yang, and Jian-chun Cheng. Acoustic one-way metasurfaces: Asymmetric phase modulation of sound by subwavelength layer. *Scientific reports*, 6(1):28023, 2016. doi: <https://doi.org/10.1038/srep28023>.
- [107] Gianluca Memoli, Mihai Caleap, Michihiro Asakawa, Deepak R Sahoo, Bruce W Drinkwater, and Sriram Subramanian. Metamaterial bricks and quantization of metasurfaces. *Nature communications*, 8(1):14608, 2017. doi: <https://doi.org/10.1038/ncomms14608>.
- [108] Yong Li, Xue Jiang, Bin Liang, Jian-chun Cheng, and Likun Zhang. Metascreen-based acoustic passive phased array. *Physical Review Applied*, 4(2):024003, 2015. doi: <https://doi.org/10.1103/PhysRevApplied.4.024003>.
- [109] Chen Shen, Yangbo Xie, Junfei Li, Steven A Cummer, and Yun Jing. Asymmetric acoustic transmission through near-zero-index and gradient-index metasurfaces. *Applied Physics Letters*, 108(22), 2016. doi: <https://doi.org/10.1063/1.4953264>.
- [110] Steven A Cummer, Johan Christensen, and Andrea Alù. Controlling sound with acoustic metamaterials. *Nature Reviews Materials*, 1(3):1–13, 2016. doi: <https://doi.org/10.1038/natrevmats.2016.1>.

- [111] Guancong Ma and Ping Sheng. Acoustic metamaterials: From local resonances to broad horizons. *Science advances*, 2(2):e1501595, 2016. doi: 10.1126/sciadv.1501595.
- [112] Mourad Oudich, Yong Li, Badreddine M Assouar, and Zhilin Hou. A sonic band gap based on the locally resonant phononic plates with stubs. *New Journal of Physics*, 12(8):083049, 2010. doi: 10.1088/1367-2630/12/8/083049.
- [113] Tai-Yun Huang, Chen Shen, and Yun Jing. Membrane-and plate-type acoustic metamaterials. *The Journal of the Acoustical Society of America*, 139(6):3240–3250, 2016. doi: <https://doi.org/10.1121/1.4950751>.
- [114] Xue Jiang, Yong Li, and Likun Zhang. Thermoviscous effects on sound transmission through a metasurface of hybrid resonances. *The Journal of the Acoustical Society of America*, 141(4):EL363–EL368, 2017. doi: <https://doi.org/10.1121/1.4979682>.
- [115] Nikhil JRK Gerard, Yong Li, and Yun Jing. Investigation of acoustic metasurfaces with constituent material properties considered. *Journal of Applied Physics*, 123(12), 2018. doi: <https://doi.org/10.1063/1.5007863>.
- [116] Yong Li, Gaokun Yu, Bin Liang, Xinye Zou, Guangyun Li, Su Cheng, and Jianchun Cheng. Three-dimensional ultrathin planar lenses by acoustic metamaterials. *Scientific reports*, 4(1):6830, 2014. doi: <https://doi.org/10.1038/srep06830>.
- [117] Yifan Zhu, Jie Hu, Xudong Fan, Jing Yang, Bin Liang, Xuefeng Zhu, and Jianchun Cheng. Fine manipulation of sound via lossy metamaterials with independent and arbitrary reflection amplitude and phase. *Nature communications*, 9(1):1632, 2018. doi: <https://doi.org/10.1038/s41467-018-04103-0>.
- [118] Zixian Liang and Jensen Li. Extreme acoustic metamaterial by coiling up space. *Physical review letters*, 108(11):114301, 2012. doi: <https://doi.org/10.1103/PhysRevLett.108.114301>.
- [119] Yong Li, Bin Liang, Xu Tao, Xue-feng Zhu, Xin-ye Zou, and Jian-chun Cheng. Acoustic focusing by coiling up space. *Applied Physics Letters*, 101(23), 2012. doi: <https://doi.org/10.1063/1.4769984>.

- [120] Tobias Frenzel, Jan David Brehm, Tiemo Bückmann, Robert Schittny, Muamer Kadic, and Martin Wegener. Three-dimensional labyrinthine acoustic metamaterials. *Applied Physics Letters*, 103(6), 2013. doi: <https://doi.org/10.1063/1.4817934>.
- [121] Zixian Liang, Tianhua Feng, Shukin Lok, Fu Liu, Kung Bo Ng, Chi Hou Chan, Jinjin Wang, Seunghoon Han, Sangyoon Lee, and Jensen Li. Space-coiling metamaterials with double negativity and conical dispersion. *Scientific reports*, 3(1):1614, 2013. doi: <https://doi.org/10.1038/srep01614>.
- [122] Yong Li, Bin Liang, Xin-ye Zou, and Jian-chun Cheng. Extraordinary acoustic transmission through ultrathin acoustic metamaterials by coiling up space. *Applied Physics Letters*, 103(6), 2013. doi: <https://doi.org/10.1063/1.4817925>.
- [123] Yangbo Xie, Bogdan-Ioan Popa, Lucian Zigoneanu, and Steven A Cummer. Measurement of a broadband negative index with space-coiling acoustic metamaterials. *Physical review letters*, 110(17):175501, 2013. doi: <https://doi.org/10.1103/PhysRevLett.110.175501>.
- [124] Jiajun Zhao, Baowen Li, Zhining Chen, and Cheng-Wei Qiu. Manipulating acoustic wavefront by inhomogeneous impedance and steerable extraordinary reflection. *Scientific reports*, 3(1):2537, 2013. doi: <https://doi.org/10.1038/srep02537>.
- [125] Jiajun Zhao, Baowen Li, Zhi Ning Chen, and Cheng-Wei Qiu. Redirection of sound waves using acoustic metasurface. *Applied Physics Letters*, 103(15), 2013. doi: <https://doi.org/10.1063/1.4824758>.
- [126] Bingyi Liu, Wenyu Zhao, and Yongyuan Jiang. Full-angle negative reflection realized by a gradient acoustic metasurface. *AIP Advances*, 6(11), 2016. doi: <https://doi.org/10.1063/1.4967430>.
- [127] Bingyi Liu, Wenyu Zhao, and Yongyuan Jiang. Apparent negative reflection with the gradient acoustic metasurface by integrating supercell periodicity into the generalized law of reflection. *Scientific reports*, 6(1):38314, 2016. doi: <https://doi.org/10.1038/srep38314>.
- [128] Miguel Molerón, Marc Serra-Garcia, and Chiara Daraio. Acoustic fresnel lenses with extraordinary transmission. *Applied Physics Letters*, 105(11), 2014. doi: <https://doi.org/10.1063/1.4896276>.

- [129] Rasha Al Jahdali and Ying Wu. High transmission acoustic focusing by impedance-matched acoustic meta-surfaces. *Applied Physics Letters*, 108(3), 2016. doi: <https://doi.org/10.1063/1.4939932>.
- [130] Yangbo Xie, Adam Konneker, Bogdan-Ioan Popa, and Steven A Cummer. Tapered labyrinthine acoustic metamaterials for broadband impedance matching. *Applied Physics Letters*, 103(20), 2013. doi: <https://doi.org/10.1063/1.4831770>.
- [131] Ye Tian, Qi Wei, Ying Cheng, and Xiaojun Liu. Acoustic holography based on composite metasurface with decoupled modulation of phase and amplitude. *Applied Physics Letters*, 110(19), 2017. doi: <https://doi.org/10.1063/1.4983282>.
- [132] Yihang Ding, Eleftherios Christos Statharas, Kui Yao, and Minghui Hong. A broadband acoustic metamaterial with impedance matching layer of gradient index. *Applied Physics Letters*, 110(24), 2017. doi: <https://doi.org/10.1063/1.4986472>.
- [133] Min Yang, Shuyu Chen, Caixing Fu, and Ping Sheng. Optimal sound-absorbing structures. *Materials Horizons*, 4(4):673–680, 2017. doi: <https://doi.org/10.1039/C7MH00129K>.
- [134] Jun Mei, Guancong Ma, Min Yang, Zhiyu Yang, Weijia Wen, and Ping Sheng. Dark acoustic metamaterials as super absorbers for low-frequency sound. *Nature communications*, 3(1):756, 2012. doi: <https://doi.org/10.1038/ncomms1758>.
- [135] Jorge P Arenas and Malcolm J Crocker. Recent trends in porous sound-absorbing materials. *Sound & vibration*, 44(7):12–18, 2010.
- [136] Xue Jiang, Bin Liang, Rui-qi Li, Xin-ye Zou, Lei-lei Yin, and Jian-chun Cheng. Ultra-broadband absorption by acoustic metamaterials. *Applied Physics Letters*, 105(24), 2014. doi: <https://doi.org/10.1063/1.4904887>.
- [137] Uno Ingard. On the theory and design of acoustic resonators. *The Journal of the acoustical society of America*, 25(6):1037–1061, 1953. doi: <https://doi.org/10.1121/1.1907235>.
- [138] M Badreddine Assouar, Matteo Senesi, Mourad Oudich, Massimo Ruzzene, and Zhilin Hou. Broadband plate-type acoustic metamaterial for low-frequency sound attenuation. *Applied Physics Letters*, 101(17), 2012. doi: <https://doi.org/10.1063/1.4764072>.

- [139] Mourad Oudich, Matteo Senesi, M Badreddine Assouar, Massimo Ruzenne, Jia-Hong Sun, Brice Vincent, Zhilin Hou, and Tsung-Tsong Wu. Experimental evidence of locally resonant sonic band gap in two-dimensional phononic stubbed plates. *Physical Review B—Condensed Matter and Materials Physics*, 84(16):165136, 2011. doi: <https://doi.org/10.1103/PhysRevB.84.165136>.
- [140] Mourad Oudich, Xiaoming Zhou, and M Badreddine Assouar. General analytical approach for sound transmission loss analysis through a thick metamaterial plate. *Journal of Applied Physics*, 116(19), 2014. doi: <https://doi.org/10.1063/1.4901997>.
- [141] Noe Jiménez, Weichun Huang, Vicent Romero-García, Vincent Pagneux, and J-P Groby. Ultra-thin metamaterial for perfect and quasi-omnidirectional sound absorption. *Applied Physics Letters*, 109(12), 2016. doi: <https://doi.org/10.1063/1.4962328>.
- [142] Deqiang Kong, Sibao Huang, Dongting Li, Chen Cai, Zhiling Zhou, Botao Liu, Guoxin Cao, Xuefeng Chen, Yong Li, and Shengchun Liu. Low-frequency multi-order acoustic absorber based on spiral metasurface. *The Journal of the Acoustical Society of America*, 150(1):12–18, 2021. doi: <https://doi.org/10.1121/10.0005134>.
- [143] Sibao Huang, Xinsheng Fang, Xu Wang, Badreddine Assouar, Qian Cheng, and Yong Li. Acoustic perfect absorbers via spiral metasurfaces with embedded apertures. *Applied Physics Letters*, 113(23), 2018. ISSN 1077-3118. doi: 10.1063/1.5063289.
- [144] Thang Tran, Feruza A Amirkulova, and Ehsan Khatami. Broadband acoustic metamaterial design via machine learning. *Journal of Theoretical and Computational Acoustics*, 30(03):2240005, 2022. doi: <https://doi.org/10.1142/S2591728522400059>.
- [145] Yihui Wang, Wei Sha, Mi Xiao, Cheng-Wei Qiu, and Liang Gao. Deep-learning-enabled intelligent design of thermal metamaterials. *Advanced Materials*, 35(33):2302387, 2023. doi: <https://doi.org/10.1002/adma.202302387>.
- [146] Jiahui Yan, Yingli Li, Guohui Yin, Gengwang Yan, and Heow Pueh Lee. Inverse-designed multifunctional metamaterials for broadband acoustic absorption and ventilation. *Composite Structures*, 362:119106, 2025. doi: <https://doi.org/10.1016/j.compstruct.2025.119106>.
- [147] Van Hai Trinh, Johann Guilleminot, Camille Perrot, and Viet Dung Vu. Learning acoustic responses from experiments: A multiscale-informed transfer learning approach.

- The Journal of the Acoustical Society of America*, 151(4):2587–2601, 2022. doi: <https://doi.org/10.1121/10.0010187>.
- [148] Chun-Teh Chen and Grace X Gu. Machine learning for composite materials. *MRs Communications*, 9(2):556–566, 2019. doi: <https://doi.org/10.1557/mrc.2019.32>.
- [149] Le Liu, Long-Xiang Xie, Weichun Huang, Xiu Juan Zhang, Ming-Hui Lu, and Yan-Feng Chen. Broadband acoustic absorbing metamaterial via deep learning approach. *Applied Physics Letters*, 120(25):251701, 2022. doi: <https://doi.org/10.1063/5.0097696>.
- [150] Ting-Wei Liu, Chun-Tat Chan, and Rih-Teng Wu. Deep-learning-based acoustic meta-material design for attenuating structure-borne noise in auditory frequency bands. *Materials*, 16(5):1879, 2023. doi: <https://doi.org/10.3390/ma16051879>.
- [151] Caglar Gurbuz, Felix Kronowetter, Christoph Dietz, Martin Eser, Jonas Schmid, and Steffen Marburg. Generative adversarial networks for the design of acoustic metamaterials. *The Journal of the Acoustical Society of America*, 149(2):1162–1174, 2021. doi: <https://doi.org/10.1121/10.0003501>.
- [152] K Mahesh, S Kumar Ranjith, and RS Mini. A deep autoencoder based approach for the inverse design of an acoustic-absorber. *Engineering with Computers*, 40(1):279–300, 2024. doi: <https://doi.org/10.1007/s00366-023-01789-9>.
- [153] Jinhong Wu, Xingxing Feng, Xuan Cai, Xufeng Huang, and Qi Zhou. A deep learning-based multi-fidelity optimization method for the design of acoustic metasurface. *Engineering with Computers*, 39(5):3421–3439, 2023.
- [154] Krupali Donda, Yifan Zhu, Aurélien Merkel, Sheng Wan, and Badreddine Assouar. Deep learning approach for designing acoustic absorbing metasurfaces with high degrees of freedom. *Extreme Mechanics Letters*, 56:101879, 2022. doi: <https://doi.org/10.1016/j.eml.2022.101879>.
- [155] Sibao Huang, Xinsheng Fang, Xu Wang, Badreddine Assouar, Qian Cheng, and Yong Li. Acoustic perfect absorbers via helmholtz resonators with embedded apertures. *The Journal of the Acoustical Society of America*, 145(1):254–262, 2019. ISSN 1520-8524. doi: 10.1121/1.5087128.

- [156] Wim Ronald Kampinga. *Viscothermal acoustics using finite elements: Analysis tools for engineers*. Phd thesis, University of Twente, 2010.
- [157] Michael R Stinson. The propagation of plane sound waves in narrow and wide circular tubes, and generalization to uniform tubes of arbitrary cross-sectional shape. *The Journal of the Acoustical Society of America*, 89(2):550–558, 1991. doi: <https://doi.org/10.1121/1.400379>.
- [158] Lari Kela. Resonant frequency of an adjustable helmholtz resonator in a hydraulic system. *Archive of Applied Mechanics*, 79(12):1115–1125, 2009. doi: <https://doi.org/10.1007/s00419-008-0279-5>.
- [159] Hongfang Chen, Yu Fu, Ling Ling, Yujin Hu, and Li Li. Design of locally resonant acoustic metamaterials with specified band gaps using multi-material topology optimization. *Materials*, 17(14):3591, 2024. doi: <https://doi.org/10.3390/ma17143591>.
- [160] Abdelkrim Khelif, Younes Achaoui, and Boujemaa Aoubiza. Locally resonant structures for low frequency surface acoustic band gap applications. In *Acoustic metamaterials: negative refraction, imaging, lensing and cloaking*, pages 43–59. Springer, 2013. doi: https://doi.org/10.1007/978-94-007-4813-2_2.
- [161] Jingwen Guo, Xin Zhang, Yi Fang, and Ziyang Jiang. A compact low-frequency sound-absorbing metasurface constructed by resonator with embedded spiral neck. *Applied Physics Letters*, 117(22), 2020. doi: <https://doi.org/10.1063/5.0031891>.
- [162] Jung-San Chen, Tzu-Huei Kuo, Wen-Yang Lo, and Cheng-Yi Wang. Low-frequency acoustic metasurface containing series-type resonators with curled necks. *Physica Scripta*, 97(8):085705, 2022. doi: [10.1088/1402-4896/ac8026](https://doi.org/10.1088/1402-4896/ac8026).
- [163] Zhenqian Xiao, Penglin Gao, Dongwei Wang, Xiao He, Yegao Qu, and Linzhi Wu. Accelerated design of low-frequency broadband sound absorber with deep learning approach. *Mechanical Systems and Signal Processing*, 211:111228, 2024. doi: <https://doi.org/10.1016/j.ymssp.2024.111228>.
- [164] Philip McCord Morse and K. Uno Ingard. *Theoretical Acoustics*. Princeton University Press, Princeton, NJ, (1986). ISBN: 978-0691024011.

- [165] ISO 10534-2:1998. Acoustics — Determination of sound absorption coefficient and impedance in impedance tubes — Part 2: Transfer-function method. European Standard, European Committee for Standardization, Brussels, Belgium, 1998.
- [166] ASTM International. Standard Test Method for Impedance and Absorption of Acoustical Materials Using a Tube, Two Microphones, and a Digital Frequency Analysis System. Technical Report ASTM E1050-19, West Conshohocken, PA, USA, 2019. DOI: 10.1520/E1050-19.
- [167] Oluwaseyi Ogun and John Kennedy. Comparison of traditional and deep learning optimisation for the design of acoustic metamaterials. In *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, volume 270, pages 1352–1362. Institute of Noise Control Engineering, 2024. doi: https://doi.org/10.3397/IN_2024_2894.
- [168] Sanjeet Kumar Singh, Om Prakash, and Shantanu Bhattacharya. Hybrid fractal acoustic metamaterials for low-frequency sound absorber based on cross mixed micro-perforated panel mounted over the fractals structure cavity. *Scientific Reports*, 12(1): 20444, 2022. doi: <https://doi.org/10.1038/s41598-022-24621-8>.
- [169] Ashif A Fathnan, Andreas E Olk, and David A Powell. Bandwidth limit and synthesis approach for single resonance ultrathin metasurfaces. *Journal of Physics D: Applied Physics*, 53(49):495304, 2020. doi: 10.1088/1361-6463/abb390.
- [170] Chi Zhang and Xinhua Hu. Three-dimensional single-port labyrinthine acoustic meta-material: Perfect absorption with large bandwidth and tunability. *Physical Review Applied*, 6(6):064025, 2016. ISSN 2331-7019. doi: 10.1103/PhysRevApplied.6.064025.
- [171] Qian Dong, Xiaolei Song, Subhrodeep Ray, and Haijun Liu. Acoustic metamaterial with air-backed diaphragm for broadband absorption: A preliminary study. In *ASME International Mechanical Engineering Congress and Exposition*, volume 84478, page V001T01A006. American Society of Mechanical Engineers, 2020. doi: 10.1115/IMECE2020-23928.
- [172] Xiaodong Cui, Vaibhava Goel, and Brian Kingsbury. Data augmentation for deep neural network acoustic modeling. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(9):1469–1477, 2015. doi: 10.1109/TASLP.2015.2438544.

- [173] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.
- [174] Patrice Y Simard, David Steinkraus, John C Platt, et al. Best practices for convolutional neural networks applied to visual document analysis. In *Icdar*, volume 3. Edinburgh, 2003.
- [175] Navdeep Jaitly and Geoffrey E Hinton. Vocal tract length perturbation (vtlp) improves speech recognition. In *Proc. ICML workshop on deep learning for audio, speech and language*, volume 117, page 21, 2013.
- [176] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- [177] Oluwaseyi Ogun, Henry Rice, and John Kennedy. Bridging model and experiment in the design and validation of a sub-wavelength acoustic metamaterial. *Journal of Applied Physics*, 137(19), 2025. doi: <https://doi.org/10.1063/5.0270361>.
- [178] Bastien Vayssette, Nicolas Saintier, Charles Brugger, Mohamed El May, and Etienne Pessard. Numerical modelling of surface roughness effect on the fatigue behavior of ti-6al-4v obtained by additive manufacturing. *International Journal of Fatigue*, 123: 180–195, 2019. doi: <https://doi.org/10.1016/j.ijfatigue.2019.02.014>.
- [179] Zhixiong Li, Ziyang Zhang, Junchuan Shi, and Dazhong Wu. Prediction of surface roughness in extrusion-based additive manufacturing with machine learning. *Robotics and Computer-Integrated Manufacturing*, 57:488–495, 2019. doi: <https://doi.org/10.1016/j.rcim.2019.01.004>.
- [180] Fátima García-Martínez, Diego Carou, Francisco de Arriba-Pérez, and Silvia García-Méndez. Toward data-driven research: preliminary study to predict surface roughness in material extrusion using previously published data with machine learning. *Rapid Prototyping Journal*, 29(8):1640–1652, 2023. doi: <https://doi.org/10.1108/RPJ-01-2023-0028>.
- [181] Dazhong Wu, Yupeng Wei, and Janis Terpenney. Surface roughness prediction in additive manufacturing using machine learning. In *International manufacturing science*

- and engineering conference*, volume 51371, page V003T02A018. American Society of Mechanical Engineers, 2018. doi: <https://doi.org/10.1115/MSEC2018-6501>.
- [182] P Kiranmayi, Taj, M Hymavathi, V Vijaya babu, PV Vinay, and Ch Himagireesh. Application of machine learning ensemble methods for prediction of surface roughness for fused deposition modeling processed parts. In *International Conference on Additive Manufacturing*, pages 369–381. Springer, 2024. doi: https://doi.org/10.1007/978-981-97-6016-9_29.
- [183] Temesgen Batu, Hirpa G Lemu, and Hailu Shimels. Application of artificial intelligence for surface roughness prediction of additively manufactured components. *Materials*, 16(18):6266, 2023. doi: <https://doi.org/10.3390/ma16186266>.
- [184] J. A. Pike, J. Klett, V. Kunc, and C. E. Duty. Machine learning in additive manufacturing: A review of learning techniques and tasks. In *Proceedings of the 33rd Annual International Solid Freeform Fabrication Symposium – An Additive Manufacturing Conference*, Austin, Texas, USA, 2022. University of Texas at Austin. doi: [10.26153/tsw/44610](https://doi.org/10.26153/tsw/44610).
- [185] Chaitanya Bakre, Mostafa Hassanian, and Cliff Lissenden. Influence of surface roughness from additive manufacturing on laser ultrasonics measurements. In *AIP conference proceedings*, volume 2102. AIP Publishing, 2019. doi: <https://doi.org/10.1063/1.5099713>.
- [186] Hiroshige Masuo, Yuzo Tanaka, Shotaro Morokoshi, Hajime Yagura, Tetsuya Uchida, Yasuhiro Yamamoto, and Yukitaka Murakami. Influence of defects, surface roughness and hip on the fatigue strength of ti-6al-4v manufactured by additive manufacturing. *International Journal of Fatigue*, 117:163–179, 2018. doi: <https://doi.org/10.1016/j.ijfatigue.2018.07.020>.
- [187] Larisa Patricia Mocanu, Costanzo Bellini, Filippo Berto, Vittorio Di Cocco, Francesco Iacoviello, and Nima Razavi. Defects in additive manufacturing and their influence on structural integrity. In *Fatigue in Additive Manufactured Metals*, pages 181–213. Elsevier, 2024. doi: <https://doi.org/10.1016/B978-0-323-91204-4.00004-6>.
- [188] MC Brennan, JS Keist, and TA Palmer. Defects in metal additive manufacturing processes, 2021.

- [189] John A Slotwinski, Edward J Garboczi, and Keith M Hebenstreit. Porosity measurements and analysis for metal additive manufacturing process control. *Journal of research of the National Institute of Standards and Technology*, 119:494, 2014. doi: <https://doi.org/10.1063/1.5099713>.
- [190] Oluwaseyi Ogun, Henry Rice, and John Kennedy. The influence of dataset generation on the performance of deep learning for amm design. *Preliminary proceedings ISMA-USD*, 2024.
- [191] Virginia Torczon. On the convergence of pattern search algorithms. *SIAM Journal on optimization*, 7(1):1–25, 1997. doi: <https://doi.org/10.1137/S1052623493250780>.
- [192] Naser Sharafkhani. A helmholtz resonator-based acoustic metamaterial for power transformer noise control. *Acoustics Australia*, 50(1):71–77, 2022. doi: <https://doi.org/10.1007/s40857-021-00256-z>.
- [193] K Mahesh, S Kumar Ranjith, and RS Mini. Inverse design of a helmholtz resonator based low-frequency acoustic absorber using deep neural network. *Journal of Applied Physics*, 129(17), 2021. doi: <https://doi.org/10.1063/5.0046582>.
- [194] Martin Fodslette Møller. A scaled conjugate gradient algorithm for fast supervised learning. *Neural networks*, 6(4):525–533, 1993. doi: [https://doi.org/10.1016/S0893-6080\(05\)80056-5](https://doi.org/10.1016/S0893-6080(05)80056-5).
- [195] Jean Boulvert, Josué Costa-Baptista, Théo Cavalieri, Maxime Perna, Edith Roland Fotsing, Vicente Romero-García, Gwenael Gabard, Annie Ross, Jacky Mardjono, and Jean-Philippe Groby. Acoustic modeling of micro-lattices obtained by additive manufacturing. *Applied Acoustics*, 164:107244, 2020. doi: <https://doi.org/10.1016/j.apacoust.2020.107244>.
- [196] Oluwaseyi Ogun and John Kennedy. Design-space segmentation of an acoustic metamaterial for interpretability and insight. *Journal of Sound and Vibration*, page 119443, 2025. doi: <https://doi.org/10.1016/j.jsv.2025.119443>.
- [197] Oluwaseyi Ogun and John Kennedy. Deep Learning-aided design of a broadband low-frequency meta-absorber for room acoustic applications. In *Proceedings of Euronoise 2025*, Malaga, Spain, 2025.

- [198] Nicolas Kaina, Fabrice Lemoult, Mathias Fink, and Geoffroy Lerosey. Negative refractive index and acoustic superlens from multiple scattering in single negative metamaterials. *Nature*, 525(7567):77–81, 2015. doi: <https://doi.org/10.1038/nature14678>.
- [199] Zheng Yang, Jun Mei, Min Yang, C T Chan, and Ping Sheng. Acoustic metamaterial panels for sound attenuation in the 50–1000 hz regime. *Applied Physics Letters*, 93(12):122904, 2008. doi: <https://doi.org/10.1063/1.3299007>.
- [200] Christoph Molnar. *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable*. Leanpub, 2022. Available online: <https://christophm.github.io/interpretable-ml-book/>.
- [201] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. doi: <https://doi.org/10.1016/j.jcp.2018.10.045>.
- [202] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018. doi: <https://doi.org/10.48550/arXiv.1802.03426>.
- [203] J. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, pages 281–297, Berkeley, California, 1967. University of California Press.
- [204] Peter J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 1987. doi: 10.1016/0377-0427(87)90125-7.
- [205] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD-96)*, pages 226–231. AAAI Press, 1996. doi: 10.5555/3001460.3001507.
- [206] Lauri Savioja and U Peter Svensson. Overview of geometrical room acoustic modeling techniques. *The Journal of the Acoustical Society of America*, 138(2):708–730, 2015. doi: <https://doi.org/10.1121/1.4926438>.

- [207] Michael Vorländer. Computer simulations in room acoustics: Concepts and uncertainties. *The Journal of the Acoustical Society of America*, 133(3):1203–1213, 2013. doi: <https://doi.org/10.1121/1.4788978>.
- [208] Lukas Aspöck and Michael Vorländer. Room geometry acquisition and processing methods for geometrical acoustics simulation models. In *Proceedings of EuroRegio 2016*, Porto, Portugal, 2016.
- [209] Alban Bassuet, Dave Rife, and Luca Dellatorre. Computational and optimization design in geometric acoustics. *Building Acoustics*, 21(1):75–85, 2014. doi: <https://doi.org/10.1260/1351-010X.21.1.75>.
- [210] D.T. Murphy, M. Beeson, S. Shelley, A. Southern, and A. Moore. Renderair – room acoustics simulation using a hybrid digital waveguide mesh. In *Proceedings of the 124th AES Convention*, Amsterdam, The Netherlands, May 2008.
- [211] Malcolm J Crocker, Frederick M Kessler, and A John Price. *Noise and Noise Control: Volume 1*. Crc Press, 2018. doi: <https://doi.org/10.1201/9781351075008>.
- [212] J. S. Mills, A. C. Petersen, and V. Salmon. Industrial noise control manual. Technical Report NIOSH Publication No. 75-183, National Institute for Occupational Safety and Health (NIOSH), 1975.
- [213] Thomas Rossing. *Springer handbook of acoustics*. Springer Science & Business Media, 2007. doi: <https://doi.org/10.1007/978-1-4939-0755-7>.
- [214] Leo L Beranek. The notebooks of wallace c. sabine. *The Journal of the Acoustical Society of America*, 61(3):629–639, 1977. doi: <https://doi.org/10.1121/1.381348>.
- [215] Aurora Magnani, Cristina Marescotti, and Francesco Pompoli. Acoustic absorption modeling of single and multiple coiled-up resonators. *Applied Acoustics*, 186:108504, 2022. doi: <https://doi.org/10.1016/j.apacoust.2021.108504>.

Appendix A

The AMM Absorber Finetuning Process

A.1 Introduction

The development of the Acoustic Metamaterial (AMM) absorber used in this study required several design iterations and refinements. The guiding principle was to achieve a structure that is sufficiently complex to support deep learning-based analysis, while remaining amenable to manufacturing and optimisation. The following sections document the progression from the initial concept to the final design, with emphasis on the rationale behind each modification. This process illustrates the balance between theoretical ideals, experimental constraints, and practical manufacturability.

A.2 Initial concept and motivation

The original inspiration was drawn from several concepts found in the metamaterial and acoustic absorber literature [9, 142, 143, 155, 215]. The design was envisioned as a spiral-based cavity resonator coupled to a perforated top plate, leveraging geometric complexity to enrich the absorber’s acoustic response.

Key motivations at this stage:

- Provide non-trivial geometry to test data-driven models.
- Maintain analytical tractability for modelling.
- Ensure the absorber could be fabricated as a single piece with standard techniques.

At this stage, the cavity was cylindrical, containing an Archimedean spiral channel, sealed with a circularly perforated top plate.

A.3 Iterative adjustments

A.3.1 Cavity geometry: from cylinder to hexagon

- **Original design:** cylindrical cavity enclosing the spiral.
- **Problem identified:** cylindrical units are inefficient to pack into arrays. Interstitial gaps lead to wasted material when tiled, and impact acoustic performance.
- **Solution:** replace the circular cavity with a hexagonal cavity, enabling seamless tessellation in planar arrays.

- **Impact:**

- Improved packing density.
- Structural stability when multiple absorbers are assembled.
- Potential improvement in acoustic performance




Figure A.1: Cavity design shapes

A.3.2 Top plate perforation: from circle to spiral geometry

- **Original design:** circular perforation on the top plate.
- **Problem identified:** Analytical model was insensitive to variations in the aperture radius.
- **Solution:** Re-design the perforation to mimic the geometry of the first turn of the Archimedean spiral.
- **Impact:**
 - Reduced acoustic leakage at glue joints.
 - Simplified analytical model due to the geometric continuity between the top plate and the included spiral
 - Aesthetically and structurally coherent design.




Figure A.2: Perforation geometries

A.3.3 Manufacturing and analytical simplification

- **Original design:** separate fabrication of cavity and top plate.
- **Problem identified:** multi-part assembly introduces tolerance errors and bonding challenges.
- **Solution:** design the entire absorber (spiral, cavity, and top plate) as a single integrated piece.
- **Impact:**
 - Simplified manufacturing (e.g., via additive manufacturing).
 - Reduced misalignment.
 - Improved repeatability in experiments.
 - Enabled analytical convenience, since the geometry is defined as one continuous structure.

A.4 Final design characteristics

The final AMM absorber design therefore comprises:

- A hexagonal cavity housing an Archimedean spiral channel.
- A spiral-patterned perforated top plate that matches the coil geometry.
- A single-piece construction for modelling simplicity.




Figure A.3: Sample 3D-printed AMM units mounted on a 199.5×199.5 [mm] square panel. The design demonstrates the efficient dense packing of hexagonal unit cells with minimal interstitial space between the unit cells. Subfigure (a) shows the external view of the fabricated sample, while subfigure (b) presents a transparent view of (a), allowing visualisation of the internal packing.

Appendix B

Numerical study

B.1 Introduction

In the low-frequency range considered in this thesis, and given the narrow regions of the absorber geometry (such as the spiral perforation and the spiral channel), the accuracy of the numerical mesh becomes particularly important. This is due to several factors:

- **Geometric resolution:** Narrow channels, such as spiral slits or perforations, are often only a few millimetres wide (i.e., a small fraction of the operating wavelength). If the mesh is too coarse, the geometry is poorly represented and the channel cross-section becomes distorted.
- **Boundary layer effects:** Acoustic and viscous boundary layers form near the channel walls. To capture these effects, fine mesh resolution is required across the narrow dimension; otherwise, energy dissipation and impedance are under-predicted.
- **Wave resolution requirement:** A minimum of 6–10 elements per wavelength is generally recommended. In narrow channels, however, additional refinement is required across the channel thickness/diameter, not just along the wavelength. A minimum of 20 mesh elements per wavelength was used in this study.
- **Numerical stability:** Insufficient mesh resolution in confined regions can introduce spurious resonances, poor convergence, and inaccurate impedance predictions.

To ensure numerical reliability, a mesh-independence study was conducted using the simplified reduced model (SRM) described in Section 4.3. This analysis was performed near the maximum frequency of interest (i.e., 500 [Hz]), as it imposes the most stringent spatial resolution requirements due to its shortest wavelength. Ensuring mesh independence at this frequency implies acceptable accuracy for the full frequency range (50–500 [Hz]).

A representative case-study geometry was investigated for mesh convergence, with parameters $L_s = 20$ [mm], $N_t = 1.8$, $p = 1.7$ [mm], $t = 0.8$ [mm], $t_p = 1.5$ [mm], and $H = 15$ [mm].

Interpretation

- f_r Convergence: The resonance frequency stabilises around 399–402 [Hz] for all mesh levels, with a maximum variation of 1.25%. Beyond the **Normal** mesh, changes in f_r are minimal.

Table B.1: Mesh convergence analysis

Mesh Level	DoF	f_r [Hz]	α_p	Time [min]
Extremely Fine	713246	402	0.9879	476.45
Fine	177626	397	0.9741	41.75
Normal	122286	399	0.9782	25.37
Coarse	87607	399	0.9838	15.15
Extremely Coarse	53194	402	0.9923	7.57

- α_p Convergence: The peak absorption value also shows minor fluctuations. Notably, **Coarse** and **Extremely coarse** meshes slightly overestimate α_p .
- Computation Time: The jump in computation time from **Fine** (42 [min]) to **Extremely fine** (476 [min]) offers negligible accuracy improvement.

Based on the above observations, the **Normal** mesh level (DoF = 122,286) provides the optimal mesh configuration, as f_r and α_p values converges within acceptable tolerance, a manageable computation time (25 [min]), and a good balance between computational efficiency and accuracy. Figure B.1 shows the plots of the analysis. The top plot shows that both values converge by the **Normal or Fine** mesh level, and the bottom plot shows that computation time explodes at the extremely fine mesh (almost 10x increase for marginal accuracy gain), a strong justification to avoid over-refinement. Thus, a **Normal(122,286 DoF)** or **Fine(122,286 DoF)** is a good trade-off.




Figure B.1: Mesh convergence analysis. The top subplot shows the variation in resonance frequency and peak absorption with increasing mesh density (measured by degrees of freedom, DoF). The bottom subplot illustrates the corresponding increase in computational time. The mesh level with **DoF = 122,286** provides a good balance between solution accuracy and computational cost.

Appendix C

Supplementary Parametric Studies

C.1 Introduction

In addition to the main design and parametric studies presented in Chapter 4.1, two supplementary investigations were conducted on the AMM absorber. Although not central to the primary tuning process, these studies provide further insight into the absorber's behaviour under variations of selected parameters and offer useful guidance for future design exploration.

C.2 Investigated Parameters

The following variables were considered:

- **Initial radius, r_o :** While this variable is not directly controlled but depends on other geometric parameters such as $(p + t)$ and H , its impact on acoustic response is worth consideration for decision making as it affects the open area ratio of the top plate
- **Thickness, tp , of the top plate:** This parameter was kept constant throughout the study as H was used instead to control the depth of the absorber.




Figure C.1: Variations in acoustic response due to changes in tp . Here, $L_s = 35$ [mm], $Nt = 3$, $p = 2.7$ [mm], $t = 0.8$ [mm], and $H = 15$ [mm], are fixed

Figure C.1 shows that absorptivity is very less sensitive to variations in tp despite the large step size of 5 [mm] used in the simulation study, changes in the responses are very minimal. This justifies setting this value constant throughout the study.

Solving for r_o in Eq.(3.1) yields a quadratic equation, which in most cases produces one positive and one negative root. Since only positive values are geometrically meaningful for manufacturing, the negative root is discarded. Table C.1 presents the variation of r_o for different values of $(p+t)$ and H . For analytical convenience, r_o was not treated as a directly controlled parameter in this thesis.

Table C.1: Changes in r_o due to changes in $(p+t)$, keeping other parameters constant, i.e., $Ls = 35$ [mm], $Nt = 3$, $t = 0.8$ [mm], $H = 15$ [mm], and $tp = 1.5$ [mm]

$(p+t)$ [mm]	r_o [mm]	f_r [Hz]	α_p
1.0000	0.4336	52.5000	0.4579
1.5000	1.0262	103.0000	0.4087
2.0000	1.3230	124.0000	0.7549
2.5000	1.5061	137.5000	0.9158
3.0000	1.6232	148.5000	0.9837
3.5000	1.6954	157.5000	1.0000
4.0000	1.7344	165.5000	0.9890
4.5000	1.7473	173.0000	0.9648
5.0000	1.7391	180.0000	0.9335
5.5000	1.7132	187.0000	0.9010
6.0000	1.6721	194.0000	0.8685
6.5000	1.6177	201.5000	0.8363
7.0000	1.5517	208.5000	0.8067
7.5000	1.4752	216.0000	0.7776
8.0000	1.3893	224.0000	0.7503
8.5000	1.2947	232.5000	0.7245
9.0000	1.1923	241.5000	0.6998
9.5000	1.0825	251.5000	0.6763
10.0000	0.9660	262.0000	0.6539

Figure C.2 shows the effect of r_o on peak absorption based on data from Table C.1. It can be observed that a nonlinear quadratic trend exists between the variables due to the quadratic relationship between p and r_o . The parametric investigation reveals a strong coupling between the aperture initial radius, the resonant frequency, and the peak absorptivity. As p increases and thus decreases r_o , the resonance is systematically shifted toward higher frequencies due to reduced cavity space. Concurrently, the absorptivity exhibits a non-monotonic dependence on r_o . At larger values of r_o , the absorption is relatively weak due to the impedance mismatch between the absorber and the background medium. As r_o is reduced, α_p increases and

reaches values close to unity. This regime corresponds to an optimal impedance-matching condition, where the absorber is able to trap and dissipate incident energy efficiently. Beyond this optimum, further reduction of r_o results in a decline of α_p , despite the continued upward shift of the resonance frequency. This deterioration indicates that the absorber geometry departs from the matching condition, thereby reducing the efficiency of energy dissipation mechanisms. The variation in p controls both the location of the resonance and the achievable absorption peak. In effect, p tunes the coupling strength between structural geometry and acoustic dissipation, highlighting the necessity of a multi-parameter optimisation strategy. The combined influence of r_o and p demonstrates that neither frequency location nor absorption strength can be maximized independently; rather, they must be co-optimised to achieve robust acoustic performance.




Figure C.2: Plot showing the peak absorption as a function of r_o and peak frequency, f_r , under variations in $(p + t)$.

Table C.2 shows the variations in r_o due to absorber thickness, H . A quadratic relationship is also observed between the variables in Fig. C.3, highlighting the complex interplay among the absorber's geometric parameters and their influence on overall performance. The 3D plot illustrates the interdependent behavior of f_r , r_o , and α_p , as H varies. As H increases,

Table C.2: Changes in r_o due to changes in $(p+t)$, keeping other parameters constant, i.e., $Ls = 35$ [mm], $Nt = 3$, $t = 0.8$ [mm], $p = 3.5$ [mm], and $tp = 1.5$ [mm]

H [mm]	r_o [mm]	f_r [Hz]	α_p
5.0000	0.0601	180.0000	0.7301
6.0000	0.2837	176.0000	0.8086
7.0000	0.4850	173.0000	0.8666
8.0000	0.6696	170.5000	0.9079
9.0000	0.8412	168.0000	0.9400
10.0000	1.0020	165.5000	0.9614
11.0000	1.1540	164.0000	0.9774
12.0000	1.2983	162.0000	0.9891
13.0000	1.4362	160.5000	0.9954
14.0000	1.5683	159.0000	0.9987
15.0000	1.6954	157.5000	1.0000
16.0000	1.8179	156.0000	0.9984
17.0000	1.9364	155.0000	0.9964
18.0000	2.0512	154.0000	0.9916
19.0000	2.1626	152.5000	0.9873
20.0000	2.2709	151.5000	0.9822
21.0000	2.3764	150.5000	0.9761
22.0000	2.4793	149.5000	0.9691
23.0000	2.5797	149.0000	0.9630
24.0000	2.6778	148.0000	0.9563
25.0000	2.7738	147.0000	0.9485
26.0000	2.8678	146.5000	0.9416
27.0000	2.9600	145.5000	0.9338
28.0000	3.0504	145.0000	0.9267
29.0000	3.1391	144.5000	0.9184
30.0000	3.2262	143.5000	0.9115

r_o increases, and f_r decreases, indicating an inverse relationship due to an increased effective acoustic spiral path length and cavity volume. α_p , indicated by the colour gradient, increases from approximately 0.7 to near unity as r_o increases and f_r decreases, suggesting enhanced acoustic impedance matching and energy dissipation at larger r_o values. Variation in H modulates this relationship by enabling control over the resonant behavior and absorption performance, where increasing H facilitates larger r_o values corresponding to lower f_r and higher α_p .




Figure C.3: Plot showing the peak absorption as a function of initial radius, r_o and peak frequency, f_r , under variations in H .

Appendix D

Linear Acoustic Wave Theory

D.1 Theory of acoustic wave propagation

The propagation of sound waves in a medium is governed by the linear acoustic wave equation, which describes the sound pressure variations. When a sound wave propagates through a fluid medium, it creates a disturbance that varies over both space and time. In a fluid medium, the total instantaneous pressure at any point is defined as:

$$P = P_0 + p \quad (\text{D.1})$$

where p_0 is the mean (equilibrium) pressure of the medium, and p is the acoustic pressure.

Assuming a closed system (constant mass), changes in pressure lead to corresponding changes in density due to alternating high- and low-pressure zones. Likewise, the medium's velocity also fluctuates because the disturbance propagates through particle oscillations. Consequently, the medium density and particle velocity are similarly expressed as:

$$\rho = \rho_0 + \rho', \quad v = v_0 + v, \quad (\text{D.2})$$

where ρ_0 is the equilibrium density and v_0 is the mean velocity. Assuming there is no mean flow, we set $v_0 = 0$.

Acoustic pressure fluctuations are extremely small, typically ranging in the order of 10^{-10} to 10^{-4} times the equilibrium pressure of the medium. Despite their minuscule magnitude, these fluctuations are perceived as sound. Since acoustic pressure variations involve only small-scale compressions and expansions, the resulting temperature gradient is negligible, and thermal conductivity effects are also disregarded. Therefore, the acoustic process is considered reversible and adiabatic (isentropic). From the theory of thermodynamics, the relationship between the pressure and the density variation of a real gas is defined by the Taylor series expansion about ρ_0 as

$$P = P_0 + \left(\frac{\partial P}{\partial \rho} \right)_{\rho_0} (\rho - \rho_0) + \frac{1}{2} \left(\frac{\partial^2 P}{\partial \rho^2} \right)_{\rho_0} (\rho - \rho_0)^2 + \dots \quad (\text{D.3})$$

For sufficiently small acoustic fluctuations, $\rho - \rho_0 \ll 1$, so that $(\rho - \rho_0)^2$ and higher powers may be neglected, the adiabatic relationship of the fluid medium through which a sound wave propagates reduces to

$$P = P_0 + \left(\frac{\partial P}{\partial \rho} \right)_{\rho_0} (\rho - \rho_0). \quad (\text{D.4})$$

The adiabatic bulk modulus, B , defined as the resistance to compression or expansion of the fluid medium, is mathematically defined as

$$B = \rho_0 \left(\frac{\partial P}{\partial \rho} \right)_{\rho_0}. \quad (\text{D.5})$$

Substituting Eq. (D.5) into Eq. (D.4), we have that

$$P = P_0 + B \frac{(\rho - \rho_0)}{\rho_0}, \quad (\text{D.6})$$

and the acoustic pressure becomes

$$p = P - P_0 = B \frac{(\rho - \rho_0)}{\rho_0}. \quad (\text{D.7})$$

Consider an infinitesimal volume element of the fluid as shown in Fig. D.1 and applying the mass conservation law to each direction, i.e., the rate of increase in mass within the volume is equal to the net mass flowing into the infinitesimal volume. For each direction, the law can be stated mathematically as

$$\begin{aligned} \left[\rho u_x - \left(\rho u_x + \frac{\partial(\rho u_x)}{\partial x} dx \right) \right] dy dz &= - \frac{\partial(\rho u_x)}{\partial x} dV. \\ \left[\rho u_y - \left(\rho u_y + \frac{\partial(\rho u_y)}{\partial y} dy \right) \right] dz dx &= - \frac{\partial(\rho u_y)}{\partial y} dV. \\ \left[\rho u_z - \left(\rho u_z + \frac{\partial(\rho u_z)}{\partial z} dz \right) \right] dx dy &= - \frac{\partial(\rho u_z)}{\partial z} dV. \end{aligned} \quad (\text{D.8})$$

Combining the three equations in Eq. (D.8), we have that




Figure D.1: Infinitesimal fluid volume element

$$-\left(\frac{\partial(\rho u_x)}{\partial x} + \frac{\partial(\rho u_y)}{\partial y} + \frac{\partial(\rho u_z)}{\partial z}\right) dV = \nabla \cdot (\rho \mathbf{u}) dV \Rightarrow \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0, \quad (\text{D.9})$$

where ∇ is the laplacian operator, and $\mathbf{u} = u_x \hat{x} + u_y \hat{y} + u_z \hat{z}$ and $\frac{\partial \rho}{\partial t}$ is the net increase in mass (or increase in density) per unit time. Consider the same fluid element in Fig. D.1, the force experienced by an infinitesimal fluid element in the x-direction on an infinitesimal area $dydz$

$$df_x = \left[P - \left(P + \frac{\partial P}{\partial x} \right) \right] dydz = -\frac{\partial P}{\partial x} dV. \quad (\text{D.10})$$

Similarly, $df_y = -\frac{\partial P}{\partial y} dV$ and $df_z = -\frac{\partial P}{\partial z} dV$, so that the total force in all direction can be written as

$$d\mathbf{f} = -\nabla P dV. \quad (\text{D.11})$$

The net acceleration of an infinitesimal fluid element is given by

$$\mathbf{a} = \frac{\partial \mathbf{u}}{\partial t} + u_x \frac{\partial \mathbf{u}}{\partial x} + u_y \frac{\partial \mathbf{u}}{\partial y} + u_z \frac{\partial \mathbf{u}}{\partial z} = \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u}. \quad (\text{D.12})$$

The Newton's second law ($d\mathbf{f} = \mathbf{a} dm$) can be written using Eqs. (D.11) and (D.12) to obtain

$$-\nabla P dV = \rho \left(\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} \right) dV. \quad (\text{D.13})$$

Assuming that $|(\mathbf{u} \cdot \nabla) \mathbf{u}| \ll \frac{\partial \mathbf{u}}{\partial t}$. Since $|(\mathbf{u} \cdot \nabla) \mathbf{u}|$ is a nonlinear term of the second order in perturbation, it can be neglected and Eq. (D.13) then reduces to

$$\rho \frac{\partial \mathbf{u}}{\partial t} = -\nabla p. \quad (\text{D.14})$$

Equation (D.14) is the linear Euler's equation. From Eq. (D.7), let S be the condensation rate, so that $S = \frac{\rho - \rho_0}{\rho_0}$ and from which

$$\rho = \rho_0(1 + S). \quad (\text{D.15})$$

Differentiating Eq. (D.15) with respect to (w.r.t) t assuming that the medium is homogeneous with constant mean density, ρ_0 , we have that $\frac{\partial \rho}{\partial t} = \rho_0 \frac{\partial S}{\partial t}$. Substituting this into Eq. (D.9) yields

$$\rho_0 \frac{\partial S}{\partial t} + \nabla \cdot (\rho u) = 0. \quad (\text{D.16})$$

Differentiating Eq. D.16 w.r.t t , we have that

$$\rho_0 \frac{\partial^2 S}{\partial t^2} + \nabla \cdot \left(\rho \frac{\partial u}{\partial t} \right) = 0. \quad (\text{D.17})$$

Differentiating Eq. (D.14) w.r.t space, we have

$$\nabla \cdot \left(\rho \frac{\partial u}{\partial t} \right) = -\nabla^2 p. \quad (\text{D.18})$$

Combining Eqs. (D.17) and (D.18) results in

$$\rho_0 \frac{\partial^2 S}{\partial t^2} = \nabla^2 p. \quad (\text{D.19})$$

In an homogeneous medium, the bulk modulus, B and ρ_0 are constant. It can be easily deduced that $p = BS$, and upon differentiating w.r.t time, we have

$$\frac{\partial^2 p}{\partial t^2} = B \frac{\partial^2 S}{\partial t^2} \Rightarrow \frac{\partial^2 S}{\partial t^2} = \frac{1}{B} \frac{\partial^2 p}{\partial t^2}. \quad (\text{D.20})$$

Combining Eqs. (D.19) and (D.20) yields

$$\nabla^2 p - \frac{\rho_0}{B} \frac{\partial^2 p}{\partial t^2} = 0. \quad (\text{D.21})$$

Equation (D.21) can be written in a familar form by letting $c = \sqrt{\frac{B}{\rho_0}}$, where c is the thermodynamic speed of sound, to get

$$\nabla^2 p - \frac{1}{c^2} \frac{\partial^2 p}{\partial t^2} = 0. \quad (\text{D.22})$$

Equation (D.22) is the fundamental equation of linear acoustics. Throughout this project, a one-dimensional harmonic plane wave propagation is assumed.

D.1.1 Solution of the acoustic wave equation

A general sound wave can be assumed to be a combination of periodic waves with a harmonic time dependency (Fourier's theorem). Considering a plane wave traveling in the z -direction, a general harmonic solution, $p(z, t)$, as a function of time and space may be defined as $p(z, t) = A \cos(kz \pm Ct)$, where k and C are constants defined as $k = 2\pi/\lambda$ and $C = \omega$, and $\pm$ denotes the wave traveling in the $-z$ direction (+) and $+z$ direction (-). The wavenumber, k , is derived by setting $t = 0$ and periodically varying z , where the wavelength, λ , is equal to a period of the wave, and C is derived by setting $z = 0$ and periodically varying t . For analytical convenience, the equation of a harmonic plane wave travelling in the z -direction can be written in exponential form as

$$p(z, t) = Ae^j(\omega t \pm kz), \quad (\text{D.23})$$

since according to Euler's formula, $e^j(\omega t \pm kz) = \cos(\omega t \pm kz) + j\sin(\omega t \pm kz)$, where only the real part is considered in practice. It is straightforward to verify that Eq. (D.23) satisfies Eq. (D.22) by taking the second derivative of $p(z, t)$ w.r.t space and time, and substituting the results into Eq. (D.22).

For $S \ll 1$, $\rho \approx \rho_0$, and from Eq. (D.14), considering wave propagation in the z -direction only, we have that

$$\rho_0 \frac{\partial \mathbf{u}}{\partial t} = -\frac{\partial p}{\partial z}. \quad (\text{D.24})$$

Differentiating $p(z, t)$ w.r.t z , substituting the result into Eq. (D.24) and subsequently integrating both sides to obtain the equation of acoustic particle velocity travelling in the z direction as

$$u(z, t) = -\frac{1}{\rho_0 c} Ae^j(\omega t \pm kz) = \pm \frac{p_{\pm}}{\rho_0 c}, \quad (\text{D.25})$$

where p_+ is the acoustic pressure of forward propagating wave, and p_- is the acoustic pressure of backward propagating wave.

D.1.2 Sound behaviour at medium boundaries

A harmonic plane wave normally incident on a planar medium boundary would have some of its energy reflected into the incident medium, some absorbed, and some transmitted into the second medium as depicted in Fig. D.2.




Figure D.2: A harmonic plane wave with incident acoustic pressure, p_i , originates from medium 1 and is normally incident on a planar boundary at $z = 0$. The wave produces a reflected pressure, p_r , in medium 1 and a transmitted pressure, p_t , in medium 2.

From Eq. (D.23), the incident pressure p_i and reflected pressure, p_r , can be defined as

$$\begin{aligned} p_i &= p_{i,max} e^{j(\omega t - k_1 z)}, \\ p_r &= p_{r,max} e^{j(\omega t + k_1 z)} = R p_{i,max} e^{j(\omega t + k_1 z)}, \end{aligned} \quad (D.26)$$

where $R = \frac{p_{r,max}}{p_{i,max}}$, is the reflection coefficient. ω is the same for both media since it is determined by the source. However, $c_1 \neq c_2$ which implies that $k_1 \neq k_2$, differs between the two media. The net acoustic pressure at the boundary is

$$p(0, t) = p_i(0, t) + p_r(0, t) = p_{i,max}(1 + R)e^{j(\omega t)}. \quad (D.27)$$

Similarly, the acoustic velocity for the forward (incident) and backward (reflected) propagating waves using Eq. (D.25) are, respectively,

$$\begin{aligned} v_i &= \frac{p_{i,max} e^{j(\omega t - k_1 z)}}{\rho_1 c_1}, \\ v_r &= -\frac{v_{r,max} e^{j(\omega t + k_1 z)}}{\rho_1 c_1} = -\frac{R p_{i,max} e^{j(\omega t + k_1 z)}}{\rho_1 c_1}. \end{aligned} \quad (D.28)$$

The net acoustic velocity at the boundary is similarly written as

$$v(0, t) = v_i(0, t) + v_r(0, t) = \frac{v_{i,max}(1 - R)e^{j\omega t}}{\rho_1 c_1}. \quad (D.29)$$

For normal incidence, the specific acoustic impedance is the same as the normal impedance. The normal impedance at the boundary, $Z_{boundary}$, is defined as

$$Z_{boundary} = \frac{p(0, t)}{v(0, t)} = \rho_1 c_1 \frac{p_{i,max}(1 + R)e^{j\omega t}}{p_{i,max}(1 - R)e^{j\omega t}}. \quad (D.30)$$

Thus, the the impedance at the boundary, or Z_2 , then becomes

$$Z_2 = \frac{1 + R}{1 - R} Z_1 \Rightarrow \frac{Z_2}{Z_1} = \frac{1 + R}{1 - R}, \quad (D.31)$$

and the reflection coefficient, R , yields

$$R = \frac{Z_2 - Z_1}{Z_2 + Z_1}. \quad (D.32)$$

The transmitted pressure, $p_t = p_{t,max} e^{j(\omega t - k_2 z)}$. Applying continuity of pressure at the boundary,

$$p_i(0, t) + p_r(0, t) = p_t(0, t) \Rightarrow p_{i,max} e^{j\omega t} + p_{r,max} e^{j\omega t} = p_{t,max} e^{j\omega t} \quad (D.33)$$

Dividing through Eq. (D.33) by $p_{i,max} e^{j\omega t}$ results in $T = 1 + R$, where T is the transmission coefficient. For an absorption problem, the effectiveness of the boundary at absorbing sound can be quantified by the absorptivity, $\alpha = 1 - |R|^2$, i.e., the fraction of the incident acoustic energy that is converted to heat within the material. In a practical measurement, R is determined based on a two-microphone transfer function method [165].