forum acusticum 2026
8-12 September · Graz, Austria

A NEURAL EIGENVALUE SOLVER USING PLANE WAVE REPRESENTATIONS

Matteo Calafà[1,3], Allan P. Engsig-Karup[2,3], Cheol-Ho Jeong[1]

[1] *Acoustic Technology, Department of Electrical and Photonics Engineering, Technical University of Denmark, 2800 Kongens Lyngby, Denmark*

[2] *Department of Applied Mathematics and Computer Science, Technical University of Denmark, 2800 Kongens Lyngby, Denmark*

[3] *Pioneer Centre for Artificial Intelligence, Øster Voldgade 3, 1350 Copenhagen, Denmark*

Abstract

Estimating acoustic eigenfrequencies is essential in the design of rooms and buildings, particularly at low frequencies and/or small rooms. While analytical solutions exist for simple geometries, more general shapes and materials require numerical algorithms. In recent years, machine learning approaches for eigenvalue problems have been proposed, but their accuracy and robustness remain limited and unexplored, leaving finite element methods as the preferred standard. Although, a neural network architecture called HergNet has been recently introduced which leverages the plane wave decomposition to achieve superior performance over previous physics-based machine learning approaches. In this work, we present a framework that employs HergNet as an eigenvalue solver, obtaining substantially reduced training times while maintaining high accuracy and robustness. The learning mechanism further allows the search to be restricted to a specified eigenfrequency range. The method is tested against FEM across different geometries and boundary conditions, demonstrating its applicability to a wide range of scenarios.

Keywords: *eigenvalue problems, physics-constrained machine learning, plane wave decomposition*

1. Introduction

Sound waves in enclosed spaces can resonate at specific frequencies, giving rise to distinct room modes, particularly in the low-frequency range [1]. The analysis of acoustic room modes is therefore essential for understanding sound field behavior and for optimizing the design of spaces such as classrooms, theaters, and offices [2].

Mathematically, room mode analysis requires solving the Laplace eigenvalue problem given some prescribed boundary conditions. Although, eigenvalues are known explicitly only in a few special configurations, such as rectangular, cylindrical and spherical enclosures with perfectly rigid boundaries [1], [3]. For general domains and boundary conditions, numerical strategies are instead required; current software libraries typically proceed by first discretizing the underlying differential eigenvalue problem in some finite-dimensional space, for example using the finite element method (FEM) [4]. The resulting matrix eigenvalue problem is then solved using iterative techniques based on the Arnoldi [5] or Lanczos [6] algorithms and their variants. Other techniques are also available: among others, the matrix pencil method [7] offers a way to estimate the eigenvalues from a time signal, whereas Calafà et al. [8] recently attempted to extend the known formula for rectangular rooms to more general boundary conditions.

Artificial neural networks (ANNs) have also been explored as eigenvalue solvers, with seminal contributions including the work of Cichocki and Unbehauen [9] and Lagaris et al. [10]. However, their slower convergence, sensitivity to hyperparameters, and other factors limit the applicability of ANNs to large-scale eigenvalue problems. In contrast, Calafà et al. [11] recently introduced a neural architecture that constructs functions satisfying the Helmholtz equation by design, achieving improved performance in forward boundary value problems in room acoustics. In this work, we investigate this architecture, termed HergNet, as an eigenvalue solver.

The rest of the manuscript is organized as follows: Section 2 introduces the HergNet framework and the modifications applied to adapt it as an eigenvalue

Corresponding author: *matcal@dtu.dk.*

solver. Three numerical tests are carried out in Section 3 and a final discussion is presented in Section 4.

2. Method

We consider a bounded and simply connected domain $\Omega \subset \mathbb{R}^d$, $d = 1, 2, 3$, representing the acoustic enclosure. Let $\beta : \partial\Omega \to \mathbb{C}$ denote the characteristic acoustic surface admittance of the boundaries. Then, under the time-harmonic condition, the sound pressure $p : \Omega \to \mathbb{C}$ at the corresponding wavenumber $k \in \mathbb{R}$ solves the Helmholtz problem

$$\begin{cases} \nabla^2 p + k^2 p = 0, & \text{in } \Omega, \\ \nabla p \cdot \mathbf{n} + ik\beta p = 0, & \text{on } \partial\Omega, \end{cases} \tag{1}$$

where ∇^2 denotes the Laplacian operator and $\mathbf{n}$ is the outward unit normal vector to $\partial\Omega$. The associated eigenvalue problem can be stated as: find $\varphi : \Omega \to \mathbb{C}$, $\hat{k} \in \mathbb{C}$ such that

$$\begin{cases} \nabla^2 \varphi + \hat{k}^2 \varphi = 0, & \text{in } \Omega, \\ \nabla \varphi \cdot \mathbf{n} + ik\beta \varphi = 0, & \text{on } \partial\Omega, \end{cases} \tag{2}$$

and $\varphi \neq 0$. By comparing Equation (1) and (2), we emphasize that $\hat{k}$ is a complex-valued unknown, in contrast with k that is real-valued and fixed.

The forward pass of HergNet [11] applies

$$\tilde{p}(\mathbf{x}) = \frac{1}{M} \sum_{j=1}^{M} e^{ik(\mathbf{x}\cdot\mathbf{s}_j + d_j)} \tilde{h}(\mathbf{s}_j), \qquad \mathbf{x} \in \Omega, \tag{3}$$

where $\tilde{p}$ is the generated function, M is a large integer number, $\|\mathbf{s}_j\| = 1$ is the j-th plane wave direction, $d_j \in \mathbb{R}$ the phase shift and $\tilde{h}(\cdot)$ the amplitude. Specifically, the method can find p in Equation (1) by learning the model parameters: $\mathbf{s}_j, d_j$ are tunable coefficients, with $\mathbf{s}_j$ restricted to the unit sphere, and $\tilde{h}$ is represented by a multilayer perceptron (MLP). Through the representation in Equation (3), the output is expressed as a superposition of plane waves, thereby automatically satisfying the Helmholtz equation. As a result, Equation (1) is enforced solely through the boundary conditions, leading to improved performance compared to traditional machine-learning-based solvers. Moreover, universal approximation properties are recovered for sufficiently large M as a consequence of the plane wave decomposition theorem.

To account for the eigenvalue problem in Equation (2), we consider instead

$$\tilde{\varphi}(\mathbf{x}) = \frac{1}{M} \sum_{j=1}^{M} e^{i\hat{k}_R(\mathbf{x}\cdot\mathbf{s}_j + d_j)} e^{-\hat{k}_I(\mathbf{x}\cdot\mathbf{s}_j)} \tilde{h}(\mathbf{s}_j), \qquad \mathbf{x} \in \Omega,$$

where, compared to Equation (3), $\hat{k}_R, \hat{k}_I$ are also learnable and represent the real and imaginary part of the eigenvalue, respectively. By incorporating an additional tensor dimension, multiple eigenvalues can be computed concurrently; let us denote with $\{\tilde{\varphi}_n\}_{n=1}^N$ the $N \in \mathbb{N}$ stacked outputs.

To solve Equation (2), we define the objective function $\mathcal{L} := \alpha_{BC}\mathcal{L}_{BC} + \alpha_0\mathcal{L}_0 + \alpha_G\mathcal{L}_C$, where $\alpha_* > 0$ are fixed hyperparameters that determine the relative contribution of each loss term. The individual components are defined as follows:

$$\mathcal{L}_{BC} := \mathbb{E}_{n=1,\dots,N}\mathbb{E}_{\mathbf{x}\sim\mathcal{U}(\partial\Omega)} \left|\nabla\tilde{\varphi}_n \cdot \mathbf{n} + ik\beta\tilde{\varphi}_n\right|^2,$$

is aimed at enforcing the boundary condition in Equation (2), where the average is computed over the training set uniformly distributed on $\partial\Omega$. Then,

$$\mathcal{L}_0 := \max_{n=1,\dots,N} e^{-\eta_0 \mathbb{E}_{\mathbf{x}\sim\mathcal{U}(\partial\Omega)}|\tilde{\varphi}_n|},$$

for $\eta_0 > 0$ is required to prevent convergence to the trivial solution $\tilde{\varphi} = 0$.

To prevent distinct solutions from collapsing onto the same mode, we introduce a third loss term $\mathcal{L}_C$. While enforcing distinct eigenvalues $\hat{k}_n \neq \hat{k}_m$ appears straightforward, it fails to account for degenerate modes. Consequently, we instead establish a criterion based on the eigenfunctions $\tilde{\varphi}_n$. Previous approaches have addressed this by enforcing the orthogonality condition $\langle\tilde{\varphi}_n, \tilde{\varphi}_m\rangle_{L^2(\Omega)} = 0$ for $n \neq m$ [12]. However, this condition is strictly specific to the eigenvalue problem in Equation (2) and requires a dense training set $\mathbf{x} \in \Omega$ to maintain the accuracy of the quadrature formula for the $L^2(\Omega)$ scalar product. To overcome these limitations, we adopt an alternative approach. Two predicted eigenfunctions $\hat{\varphi}_n$, $\hat{\varphi}_m$ represent the identical mode if and only if they differ by a constant scaling factor $\hat{\varphi}_n = \Lambda\hat{\varphi}_m$, $\Lambda \in \mathbb{C}$. This condition holds true if $S_C(\tilde{\varphi}_n, \tilde{\varphi}_n)| = 1$, where S_C denotes the cosine similarity

$$S_C(\tilde{\varphi}_n, \tilde{\varphi}_m) := \frac{\langle\tilde{\varphi}_n, \tilde{\varphi}_m\rangle_{L^2(\Omega)}}{\|\tilde{\varphi}_n\|_{L^2(\Omega)}\|\tilde{\varphi}_m\|_{L^2(\Omega)}}.$$

Therefore, we propose the following definition:

$$\mathcal{L}_C := \max_{n>m} \sup_{\mathbf{x}\sim\mathcal{U}(\Omega)} e^{-\eta_C(1 - |S_C(\tilde{\varphi}_n, \tilde{\varphi}_m)|)}. \tag{4}$$

Compared to the original orthogonality criterion, this formulation penalizes high similarities without rigidly enforcing a value of zero, thereby accommodating non-negligible quadrature errors. Consequently, $\mathbf{x}$ can be sparsely sampled within the domain Ω.

Enforcing distinct eigenfunctions via a mutual penalty may cause optimization competition, stalling the convergence of both terms. To mitigate this issue, we aim to penalize only one term in each pairwise comparison. Technically, this is achieved by blocking the automatic differentiation (AD) gradient for either $\tilde{\varphi}_n$ or $\tilde{\varphi}_m$ in Equation (4). Standard machine learning libraries natively support this operation; for example, JAX provides the `jax.lax.stop_gradient` function[1].

[1] `https://docs.jax.dev/en/latest/_autosummary/jax.`

The overall method is highly sensitive to the initialization of $\hat{k}_R$ and $\hat{k}_I$, as the optimization converges to eigenvalues closest to their initial values. This local convergence behavior can be leveraged to restrict the search to a specific interval. In the absence of a priori knowledge regarding the expected solutions, $\hat{k}_R$ can be sampled along the real line according to the theoretical modal density [1]. Conversely, $\hat{k}_I$ is initialized to small values or zero to ensure numerical stability.

3. Numerical tests

Numerical tests are carried out in this section. The code is implemented in `JAX 0.9.2` and run on a `NVIDIA H100 PCIe` GPU, `80 GB VRAM`. In all tests, the speed of sound is set to $c = 343$ m/s. Furthermore, the number of plane waves in HergNet is set to $M = 200$ and the MLP for $\tilde{h}$ is composed by two hidden layers with 10 units and ReLU activation functions. Training is performed for 5000 epochs using vanilla stochastic gradient descent (SGD) for $\hat{k}_R$ and $\hat{k}_I$ with a learning rate of 0.1, while the HergNet parameters are optimized using the Adam optimizer [13] with an initial learning rate of 0.01. For convenience, we address the eigenfrequency $\hat{f}$ rather than the wavenumber, where the latter can be recovered via $\hat{f} = \hat{k}c/(2\pi)$.

3.1. Rectangular room with rigid walls

As a first simple test we consider a rectangular domain $\Omega = (0, L_x) \times (0, L_y)$, with perfectly rigid walls $\beta = 0$. The analytical solution is well-known [3] and equal to

$$\hat{f}_{n_x,n_y} = \frac{c}{2}\sqrt{\frac{n_x^2}{L_x^2} + \frac{n_y^2}{L_y^2}},$$

for the multi-index $n_x, n_y = 0, 1, \ldots$.
We consider $L_x = 1$ m, $L_y = 1.4$ m, according to the Louden's ratio [14]. $\hat{f}$ is initialized as $N = 5$ evenly spaced values from 150 Hz to 300 Hz. Specifically, N is chosen as slightly larger than the expected number of eigenvalues from the theoretical modal density. The parameter $\hat{k}_I$ is fixed to zero, as described in Section 2. 100 training points are uniformly sampled on the boundary of the domain, corresponding to approximately 23 points per wavelength (PPW) at the highest frequency of 300 Hz. An equal number of points is uniformly sampled in Ω for evaluating $\mathcal{L}_C$. The loss hyperparameters are first set to $\alpha_{BC}, \alpha_0, \eta_0 = 1$ and $\alpha_C = 0$. A second experiment is carried out with $\alpha_C = 1, \eta_C = 5$ in order to enforce the similarity penalization in Equation 4. Convergence to the eigenvalue solutions is shown in Figure 1.
Both experiments converge within a few epochs but exhibit distinct behaviors. Without the similarity penalization, convergence is fast and uniform across the N outputs; however, multiple outputs may collapse to the same solution, also resulting in missed modes, such as the solution around 250 Hz. In contrast, including

`lax.stop_gradient.html`

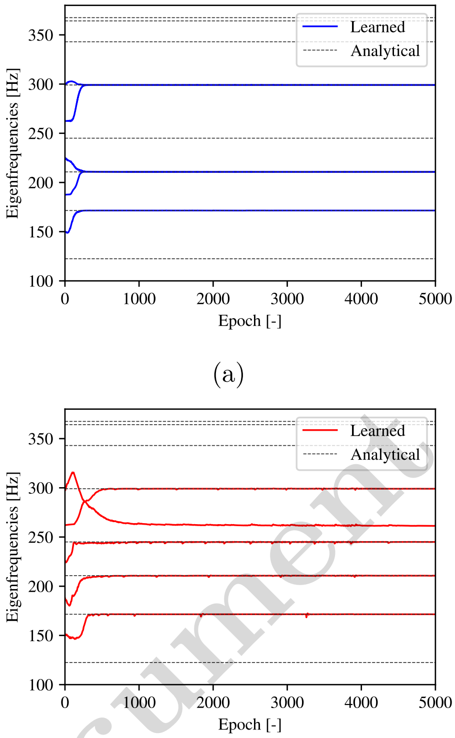

(a) (b)

Figure 1: Convergence to the eigenvalue solutions for the rectangular room with perfectly rigid walls. (a) Without similarity penalization. (b) With similarity penalization.

$\mathcal{L}_C$ ensures that each solution is captured by at most one output and therefore increasing the likelihood of recovering all distinct solutions. This comes at the cost of slightly slower convergence, and one output failing to converge due to the absence of remaining valid solutions in the neighborhood. Consequently, incorporating the similarity penalization requires a post-processing pruning step to filter out incorrect solutions by evaluating $\mathcal{L}_{BC}$ individually. In this experiment, $\mathcal{L}_{BC}$ is below 10^{-3} for all valid outputs, while the non-convergent output yields $\mathcal{L}_{BC} \approx 0.40$.
The obtained eigenfunction is compared to the analytical solution for $n_x = 1, n_y = 2$ in Figure 2, demonstrating excellent agreement.
To quantify accuracy, we evaluate the mean relative error

$$\mathcal{E} := \mathbb{E}_{n=1,\ldots N} \min_{\hat{f} \in \mathcal{F}} \sqrt{|\hat{f}_n - \hat{f}|^2/|\hat{f}|^2},$$

where we denote by $\mathcal{F}$ the set of reference solutions. This metric is computed for both experiments. When $\mathcal{L}_C$ is not included, the error is very low, with $\mathcal{E} = 5.2 \cdot 10^{-6}$. In the second experiment, $\mathcal{E} = 1.0 \cdot 10^{-4}$ by excluding the non-converged output.

3.2. Rectangular room with soft walls

A second test is carried out with the same geometry as in Section 3.1, while instead considering non-rigid

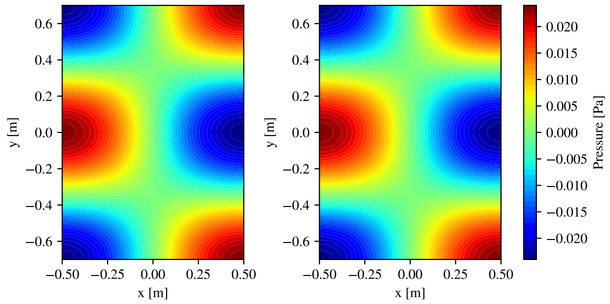

Figure 2: Learned eigenfunction (left) and analytical solution for $n_x = 1, n_y = 2$ (right) with perfectly rigid boundaries after normalization. The same color map range is used for both plots.

walls $\beta \neq 0$. In particular, we define $\zeta := 1/\beta = 10 - i$ the uniform characteristic surface impedance on the boundaries. Differently from the previous section, $\hat{f}$ is initialized as $N = 50$ evenly spaced values from 100 Hz to 600 Hz. $\mathcal{L}_C$ is for simplicity not included. The expected solutions are complex-valued due to the impedance, therefore $\hat{k}_I$ is trainable and initialized to zero. The reference solutions are obtained from the semi-analytical method proposed in Calafà et al. [8]. Convergence to the problem eigenvalues is depicted in Figure 3.

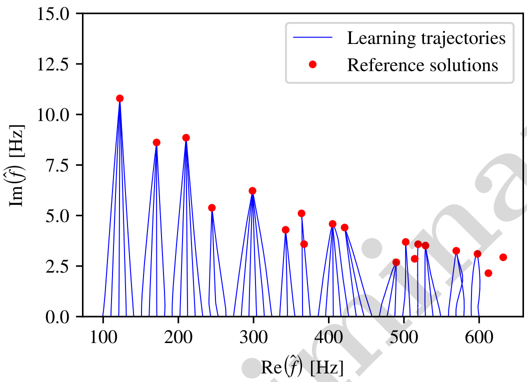

Figure 3: Trajectories of the learned eigenfrequencies during training for the rectangular room with uniform surface impedance $\zeta = 10 - i$.

Smooth convergence is observed also for this test in the complex plane. Furthermore, accurate approximations are achieved, with $\mathcal{E} = 3.0 \cdot 10^{-5}$.

3.3. L-shaped room

In this last test, we consider a non-rectangular room, with dimensions reported in Figure 4a.

Walls are considered rigid ($\beta = 0$), restricting the search to only real-valued eigenvalues. $\hat{f}$ is initialized as $N = 5$ equally spaced frequencies between 100 Hz and 350 Hz. The reference solutions are computed using the FEM-based eigenvalue solver from the Acoustics Module in COMSOL Multiphysics® [15], employing a mesh consisting of 479 quadratic Lagrangian elements. Convergence is shown in Figure 4b, which yields a final error of $\mathcal{E} = 1.4 \cdot 10^{-3}$. Specifically, the presence of the singularity induced by the concave square angle causes diffraction, leading to marginally slower convergence and increased error relative to the previous tests.

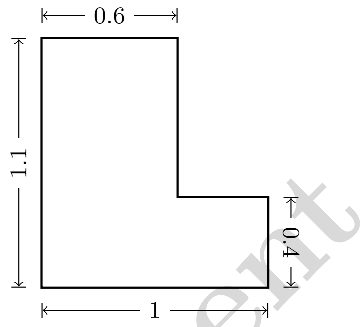

(a)

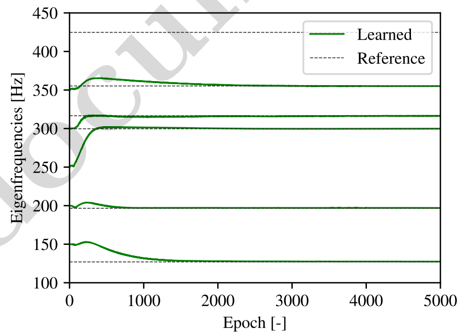

(b)

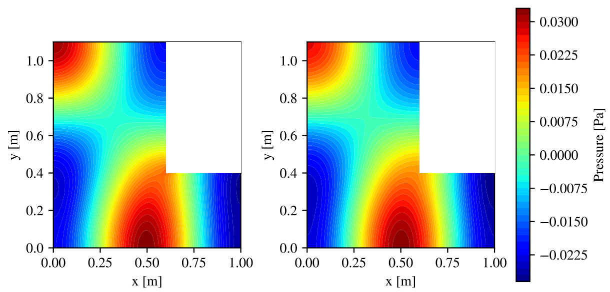

(c)

Figure 4: Details of the L-shaped room test with perfectly rigid boundaries. (a) Geometry dimensions. (b) Eigenvalue convergence during training. (c) Normalized eigenfunctions at $\hat{f} \approx 350$ Hz. Learned (left), numerical (right).

Furthermore, the contours of the learned and numerical eigenfunctions are presented in Figure 4c, where close agreement is observed once again.

4. Discussion

In this work, we demonstrated that the HergNet framework can be successfully adapted to solve eigenvalue problems. The approach achieves high accuracy, as demonstrated by the very low relative errors. In addition, it maintains high performance, requiring few training points and parameters. For example, for the test in Section 3.1, training requires approximately 8 seconds and 1.5 GB VRAM. Although recent studies on ANNs for eigenvalue problems rarely report explicit computational costs, most are expected to be more expensive due to their larger models and/or higher number of epochs [12], [16]. Furthermore, the cost of the proposed approach is only moderately higher than that of FEM, which takes 4 seconds and 1.2 GB of CPU memory for the test in Section 3.1.

On the other hand, one primary challenge lies in ensuring convergence to distinct eigenvalues without multiple solutions collapsing onto the same one. While this is a common hurdle for scientific machine learning approaches to eigenvalue problems, we proposed an effective solution that incurs negligible additional computational cost. However, this adjustment still hinders the convergence rate, particularly when using several stacked outputs. Conversely, omitting the similarity penalty typically accelerates convergence, and a large output number N often suffices to capture all eigenvalues, as demonstrated in Section 3.2. In conclusion, both approaches are valid, and the choice between them ultimately depends on the tradeoff between memory efficiency and accuracy.

Acknowledgments

Computing resources were provided by the Pioneer Centre for Artificial Intelligence, Denmark.

References

[1] H. Kuttruff, *Room acoustics*. Boca Raton, Florida: CRC Press, 2016.

[2] M. Vorländer, *Auralization: fundamentals of acoustics, modelling, simulation, algorithms and acoustic virtual reality*. Springer, 2008.

[3] P. M. Morse and K. U. Ingard, *Theoretical acoustics* (International series in pure and applied physics), eng. New York, New York: McGraw-Hill, 1968.

[4] I. Harari, "A survey of finite element methods for time-harmonic acoustics," *Computer methods in applied mechanics and engineering*, vol. 195, no. 13-16, pp. 1594–1607, 2006.

[5] W. E. Arnoldi, "The principle of minimized iterations in the solution of the matrix eigenvalue problem," *Quarterly of applied mathematics*, vol. 9, no. 1, pp. 17–29, 1951.

[6] C. Lanczos, "An iteration method for the solution of the eigenvalue problem of linear differential and integral operators," *Journal of research of the National Bureau of Standards*, vol. 45, no. 4, pp. 255–282, 1950.

[7] A. G. Prinn and E. A. Habets, "Acoustic eigenvalue estimation using Rayleigh quotient iteration with a sparse matrix pencil," *Applied Acoustics*, vol. 232, p. 110 573, 2025.

[8] M. Calafà, Y. Xia, J. Brunskog, and C.-H. Jeong, "Evaluation of acoustic Green's function in rectangular rooms with general surface impedance walls," *arXiv preprint*, 2026.

[9] A. Cichocki and R. Unbehauen, "Neural networks for computing eigenvalues and eigenvectors," *Biological Cybernetics*, vol. 68, no. 2, pp. 155–164, 1992.

[10] I. E. Lagaris, A. Likas, and D. I. Fotiadis, "Artificial neural network methods in quantum mechanics," *Computer physics communications*, vol. 104, no. 1-3, pp. 1–14, 1997.

[11] M. Calafà, Y. Xia, and C.-H. Jeong, "HergNet: A fast neural surrogate model for sound field predictions via superposition of plane waves," in *ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2026, pp. 15 197–15 201.

[12] C. Rowan, A. Doostan, K. Maute, and J. Evans, "Solving engineering eigenvalue problems with neural networks using the rayleigh quotient," *International Journal for Numerical Methods in Engineering*, vol. 126, no. 24, 2025.

[13] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint*, 2014.

[14] M. Louden, "Dimension-ratios of rectangular rooms with good distribution of eigentones," *Acustica*, vol. 24, pp. 101–104, 1971.

[15] "COMSOL Multiphysics®," version 6.4, [Online]. Available: `www.comsol.com`.

[16] I. Ben-Shaul, L. Bar, D. Fishelov, and N. Sochen, "Deep learning solution of the eigenvalue problem for differential operators," *Neural Computation*, vol. 35, no. 6, pp. 1100–1134, 2023.