Neural Operators as Fast Surrogate Models for the Transmission Loss of Parameterized Sonic Crystals

Jakob E. Wagner[* 1,2] Samuel Burbulla[* 1] Johannes D. Schmid[2]
Miguel de Benito Delgado[1]

[1]appliedAI Institute for Europe [2]TUM

Abstract

Neural operators serve as efficient, data-driven surrogate models for complex physical and engineering problems. In this work, we demonstrate that neural operators can directly learn the key properties of sonic crystals, a type of acoustic metamaterial consisting of a lattice of parameterized shapes. We predict the transmission loss curve, a critical characteristic in applications, bypassing the expensive meshing and solving steps typical of classical techniques. We evaluate established architectures, DeepONet (DON) and Fourier Neural Operator (FNO), alongside two new ones, Deep Neural Operator (DNO) and Deep Cat Operator (DCO), which demonstrate significant performance improvements. In our experiments, all models achieve high accuracy, while being up to 10^6 times faster than the traditional method, significantly advancing practical real-time metamaterial design.

1 Introduction

Metamaterials are composite materials with behavior not found in nature, whose properties arise from their internal structure rather than their composition [1]. They can be designed to have, e.g., new elastic, electromagnetic, or thermal properties [2, 3, 4], and have numerous applications including advanced sensors, antenna design, energy harvesting, and civil engineering [5, 6, 7, 8], to name a few. In acoustics, they are mostly concerned with sound attenuation [9, 10], wave manipulation [11, 12, 13], or architectural acoustics [7]. Efficient and accurate simulation methods are essential in order to predict the performance of acoustic metamaterials *in silico*. State-of-the-art approaches [14, 15, 16, 17, 18], while sufficient for this purpose, can be computationally expensive, requiring a mesh of the domain and solving optimization problems for a large number of parameters.

Sonic crystals are acoustic metamaterials composed of identical unit cells organized in a lattice, cf. Figure 1. The composition creates so-called band-gaps in the transmission loss graph, where the transmission loss (TL) describes the decrease in power from the incident to the transmitted wave [9, 19]. The graph of the TL over the frequency is a key characteristic, describing the interaction of the crystal with the environment, cf. Figure 1 (right). The primary objective in designing these metamaterials is to predict and achieve a sonic crystal with a specific TL profile.

Neural operators (NOs) [20] and related techniques [21, 22, 23, 24] have been successfully applied to model complex physical systems, particularly to approximate the solution operators of partial differential equations. They are accurate enough to complement or even replace traditional numerical simulators in areas like computational fluid dynamics, weather forecasting, and material modeling, while being orders of magnitude faster [25, 26].

In this work, we train NOs to learn the mapping from the parameterization of sonic crystals to their corresponding *transmission loss graph*, see Section 2. This bypasses the expensive step required in conventional approaches, where one must first compute the pressure solution of the system before

[*]These authors contributed equally to this work. Correspondence to: mail@jakob-wagner.org.

D3S3: Data-driven and Differentiable Simulations, Surrogates, and Solvers @ NeurIPS 2024.

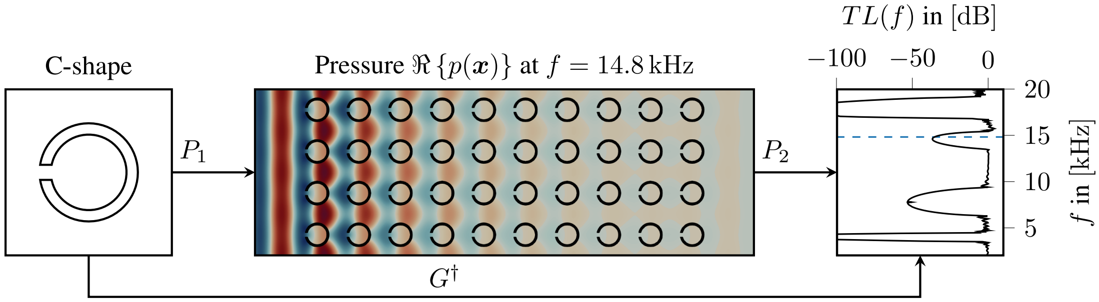

Figure 1: Our sonic crystals consist of parameterized C-shapes (left), 10 in a row in x-direction with periodicity in y-direction. The sound wave is incident from the left-hand side, traveling to the right-hand side. We learn the operator $G^\dagger$ mapping the parameterization to the transmission loss graph of the sonic crystal (right). The plot shows the sonic crystal for $R_1 = 6.5\,\mathrm{mm}$, $R_2 = 5.0\,\mathrm{mm}$, and $b = 2.0\,\mathrm{mm}$ (see text), along with the real part of the corresponding pressure field for $f = 14.8$ kHz.

evaluating the transmission loss, cf. Figure 1. We compare standard NO architectures like DeepONet (DON) [22] and FNO [21] with two new architectures: Deep Neural Operator (DNO) and Deep Cat Operator (DCO), see Section 3. All trained NOs show good generalization with slightly differing performance, see Section 4, and offer significant speed-ups over traditional methods.

2 Background and related work

Acoustics is the science that describes the behavior of sound in media, and the Helmholtz equation

$$\Delta p(x) + k^2 p(x) = 0, \tag{1}$$

is the standard modeling approach for the complex-valued sound pressure p, where $k = 2\pi f/c$ is a given *wave number* depending on a frequency f and the speed of sound c [27]. The Helmholtz equation allows for the analysis of how sound behaves under steady-state conditions, essential for understanding resonance, transmission, and the acoustic properties of materials [19].

Sonic crystals are acoustic metamaterials consisting of identical unit cells arranged in a lattice [9]. A common parameterization of sonic crystals is given by C-shapes [10], where the parameters are outer radius R_1, inner radius R_2, and opening width b of the C-shaped inclusion, cf. Figure 1 (left). A typical design objective is to prevent specific frequency ranges from propagating through the material, a phenomenon seen as *band-gaps* in the graph of the *transmission loss* [9].

The **transmission loss** [19] describes the decrease from the power W_i of the wave incident on one side of the material to the power W_t of the transmitted wave as it reaches the opposite side, $\mathrm{TL} = 10\log_{10}(W_i/W_t)$. As the transmission loss depends on the frequency f of the wave, the **transmission loss graph** $\mathrm{TL}(f)$ is the function that maps frequency to transmission loss, and is a key characteristic of sonic crystals, cf. Figure 1 (right) [9, 10].

Conventional methods like the Finite Element Method (FEM) in frequency-domain, or the Boundary Element Method (BEM) [27] are typically used to compute transmission losses. Both require a mesh of the domain for each C-shape geometry, and compute the complex-valued pressure field by solving (1) for a range of frequencies. Afterward, the resulting pressure fields can be integrated to compute the transmission loss graph [19, 28, 10].

Operator learning uses neural networks (NNs) to approximate maps between infinite dimensional function spaces, in contrast with the usual approximation of functions between Euclidean spaces. Expressing physical quantities as functions in spaces $\mathcal{A}$ and $\mathcal{U}$, operator learning approximates a physical process described by an operator $G^\dagger : \mathcal{A} \to \mathcal{U}$ with a parametric map G_θ, such that

$$G^\dagger \approx G_\theta : \mathcal{A} \to \mathcal{U}, \quad \theta \in \mathbb{R}^l, \tag{2}$$

for some number of parameters $l \in \mathbb{N}$. NOs are a flexible framework well-suited for many real-world problems. Because they are discretization-invariant, they can train and predict at different resolutions, while being much faster than traditional solvers [20].

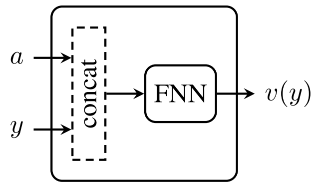

Figure 2: DNO concatenates the input function evaluations a and the evaluation location y before passing them through a single feedforward NN.

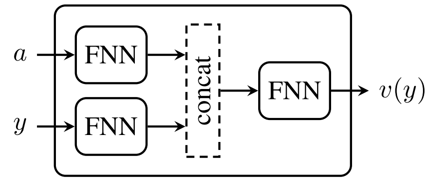

Figure 3: DCO passes the input function evaluations a and the evaluation location y through a branch and a trunk network, respectively, before passing the concatenated outputs into a third feedforward *cat network*.

DeepONet (DON, [22]) is an NO architecture motivated by the universal approximation theorem for operators. It consists of a *trunk network* T that learns a set of basis functions, and a *branch network* B that learns the corresponding basis coefficients. Given a vector a of evaluations of the input function and a coordinate y where the output function will be evaluated, a dot product linearly combines the trunk network's basis functions with the branch network's coefficients:

$$\mathcal{G}_{\text{DON}}(a)(y) = B(a) \cdot T(y). \tag{3}$$

The **Fourier Neural Operator** (FNO, [21]) has been successfully applied to various PDE-based problems with significant speedups [25]. It is a generic NO:

$$\mathcal{G}_{\text{FNO}}(a) = Q \circ (W_L + \mathcal{K}_L) \circ \cdots \circ \sigma(W_1 + \mathcal{K}_1) \circ P(a), \tag{4}$$

where P and Q are pointwise lifting and projection operators, and the intermediate layers consist of pointwise operators W_l, integral kernel operators $\mathcal{K}_l$, and an activation function σ [20]. The integral kernel operators of the FNO are linear transformations in Fourier space, $\mathcal{K}_l(v) = \mathcal{F}^{-1}(R_l \cdot \mathcal{F}(v))$, where R_l are complex-valued matrices, and $\mathcal{F}$ (resp., $\mathcal{F}^{-1}$) denotes the (inverse) Fourier transform.

3 Methodology

We consider parameterized sonic crystals and use NOs to predict their transmission loss graphs. Figure 1 shows the C-shaped unit cell we consider, and the transmission loss for a specific parameterization $a = [R_1, R_2, b]$ of the cell. We want to approximate the operator

$$G^\dagger(a)(f) = P_2\big(P_1(a, f)\big) : a \mapsto \text{TL}_a(f), \quad a \in \mathcal{A}, \quad f \in \mathbb{R}, \tag{5}$$

where TL_a is the transmission loss graph, evaluated at a frequency f, and a is understood as a constant input function. The intermediate maps P_1 and P_2 are the two parts of conventional techniques, where P_1 solves the Helmholtz equation for a given geometry and frequency, and P_2 computes the transmission loss from the resulting pressure field by integration. When training an NO, we can skip P_1 and P_2 and learn the mapping from the geometry parameters to the transmission loss graph directly.

We train a total of four different architectures: DON, FNO and our new DNO, DCO.

Deep Neural Operator (DNO) is inspired by the DeepONet architecture, although branch and trunk are replaced by a single (fully connected) neural network N processing both the input function and the evaluation coordinate simultaneously, cf. Figure 2. Using the same notation as in (3), the DNO is:

$$\mathcal{G}_{\text{DNO}}(a)(y) = N\big([a, y]\big), \tag{6}$$

where $[\cdot]$ denotes the concatenation of vectors. The motivation for DNO is to integrate operator information early by connecting all layers of the branch and trunk networks, maximizing flexibility in learning latent representations at the cost of increased computational complexity. Additionally, when the size of the concatenated tensors differ significantly, the model may disproportionately emphasize one over the other, leading to inefficiencies.

Deep Cat Operator (DCO) follows a similar motivation, but first passes the input function and the evaluation locations through branch and trunk networks, B and T, and concatenates those outputs in latent space before passing the result through a third *cat network* C, cf. Figure 3:

$$\mathcal{G}_{\text{DCO}}(a)(y) = C\big([B(a), T(y)]\big). \tag{7}$$

Essentially, DCO replaces the dot product in DeepONet by an NN, offering more flexibility for non-linear combinations of the trunk network's outputs while being more robust to disparate input tensor sizes than DNO.

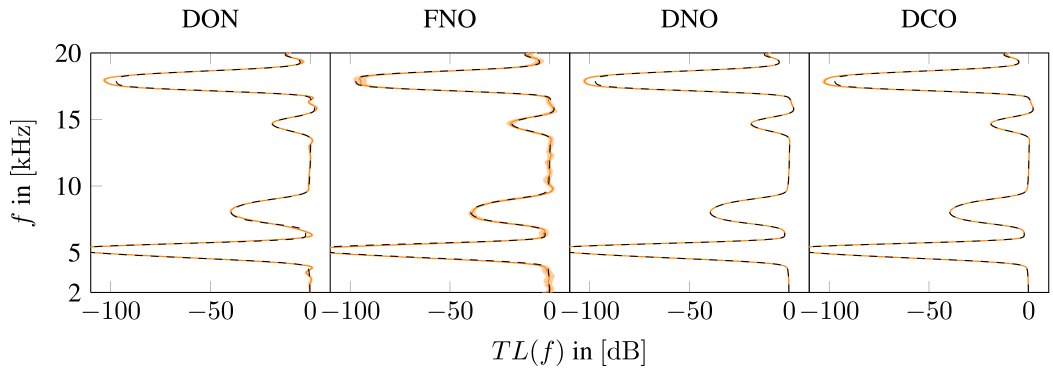

Figure 4: Mean learned transmission loss graph of four neural operators for a sample from the test set with the parameters $a = [5.7\,\mathrm{mm}, 4.5\,\mathrm{mm}, 3.2\,\mathrm{mm}]$. The darker line is the mean performance, the lighter area shows the bootstrapped 95% confidence interval with mean statistic. The dashed line is the ground truth (FEM) solution.

4 Experiments

For our experiments, we employ FEniCS [29] to produce the data set and train the NO architectures implemented in `continuiti`[2]. The setup matches the implementation in the COMSOL application gallery [28], and our dataset consists of C-shaped sonic crystals with outer radius $R_O \leq 8.5\,\mathrm{mm}$, inner radius $R_I \leq 8.5\,\mathrm{mm}$, and gap width $b \leq 8.5\,\mathrm{mm}$ embedded in a unit cell with dimensions $d = 22\,\mathrm{mm}$, while ensuring the configurations are geometrically valid. The meshes employ at least 15 elements per wavelength at $20\,\mathrm{kHz}$, and we provide 673 different parameterizations with 256 TL evaluations each for training. We train the NOs for 500 epochs using Adam to minimize the MSE, and compare the final solution to a test set with 128 parameterizations, again with 256 TL evaluations. The reported results are averaged over 10 seeds, where the runtime baseline is computed using the COMSOL implementation [28]. More details on the implementation and some additional plots can be found in Appendix A.

The plot in Figure 4 shows the transmission loss for a representative sample from the test set for all four architectures. All models show good approximation properties, with DNO and DCO yielding the best generalization properties, see Table 1.

Table 1: Performance metrics of the DON, FNO, DNO, and DCO on the test dataset obtained from $n = 10$ different initial seeds. The relative mean squared error (RMSE), the bootstrapped $95\,\%$ confidence interval (CI) for the model RMSE with mean statistic, number of parameters and speedup of all reported neural operators in comparison to our baselines.

Method	**RMSE**	$95\,\%$ **CI**	**Num. Param.**	**Speedup**
DON	$0.469\,\%$	$[0.450\,\%, 0.492\,\%]$	101 152	6.0×10^5
FNO	$2.18\,\%$	$[1.99\,\%, 2.39\,\%]$	99 553	5.7×10^6
DNO	$0.199\,\%$	$[0.196\,\%, 0.202\,\%]$	102 929	5.1×10^5
DCO	$0.208\,\%$	$[0.203\,\%, 0.212\,\%]$	100 773	1.4×10^5

5 Conclusion and future work

Our work demonstrates that neural operators can effectively learn complex relationships that require extensive computation with classical methods. The results emphasize the significance of selecting the appropriate architecture, with our new DNO and DCO models outperforming standard methods. We acknowledge the limited scope of this paper, as we focus on a single problem with a specific dataset and do not compare our approach to classical surrogate models. Thus, further experimentation is warranted, particularly to evaluate the performance of our new architectures, DNO and DCO, across a broader range of operator learning tasks. Additionally, future research could focus on directly learning the inverse problem, proposing metamaterials that achieve specific transmission loss profiles.

[2]Available on `https://github.com/aai-institute/continuiti`.

6 Acknowledgements

The appliedAI Institute for Europe gGmbH is supported by the KI-Stiftung Heilbronn gGmbH.

References

[1] Nader Engheta and Richard W Ziolkowski. *Metamaterials: physics and engineering explorations*. John Wiley & Sons, 2006.

[2] Alexandra Ion, Johannes Frohnhofen, Ludwig Wall, Robert Kovacs, Mirela Alistar, Jack Lindsay, Pedro Lopes, Hsiang-Ting Chen, and Patrick Baudisch. Metamaterial mechanisms. In *Proceedings of the 29th annual symposium on user interface software and technology*, pages 529–539, 2016.

[3] Richard A Shelby, DR Smith, SC Nemat-Nasser, and Sheldon Schultz. Microwave transmission through a two-dimensional, isotropic, left-handed metamaterial. *Applied Physics Letters*, 78(4):489–491, 2001.

[4] CZ Fan, Y Gao, and JP Huang. Shaped graded materials with an apparent negative thermal conductivity. *Applied Physics Letters*, 92(25), 2008.

[5] Yuandan Dong and Tatsuo Itoh. Metamaterial-based antennas. *Proceedings of the IEEE*, 100(7):2271–2285, 2012.

[6] Weijin Wang, Kaixiang Sun, Ying Xue, Jie Lin, Jiukai Fang, Shengnan Shi, Shan Zhang, and Yanpeng Shi. A review of terahertz metamaterial sensors and their applications. *Optics Communications*, 556:130266, 2024.

[7] Bao Zhao, Henrik R Thomsen, Jacopo M De Ponti, Emanuele Riva, Bart Van Damme, Andrea Bergamini, Eleni Chatzi, and Andrea Colombi. A graded metamaterial for broadband and high-capability piezoelectric energy harvesting. *Energy Conversion and Management*, 269:116056, 2022.

[8] Nicolás Contreras, Xihong Zhang, Hong Hao, and Francisco Hernández. Application of elastic metamaterials/meta-structures in civil engineering: A review. *Composite Structures*, 327:117663, 2024.

[9] 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.

[10] Daniel P Elford, Luke Chalmers, Feodor V Kusmartsev, and Gerry M Swallowe. Matryoshka locally resonant sonic crystal. *The Journal of the Acoustical Society of America*, 130(5):2746–2755, 2011.

[11] Yiqun Ding, Zhengyou Liu, Chunyin Qiu, and Jing Shi. Metamaterial with simultaneously negative bulk modulus and mass density. *Physical Review Letters*, 99(9):093904, 2007.

[12] Shu Zhang, Leilei Yin, and Nicholas Fang. Focusing ultrasound with an acoustic metamaterial network. *Physical Review Letters*, 102(19):194301, 2009.

[13] Sébastien Guenneau, Alexander Movchan, Gunnar Pétursson, and S Anantha Ramakrishna. Acoustic metamaterials for sound focusing and confinement. *New Journal of Physics*, 9(11):399, 2007.

[14] Yan Du, Weiguo Wu, Wei Chen, Yongshui Lin, and Qingjia Chi. Control the structure to optimize the performance of sound absorption of acoustic metamaterial: A review. *AIP Advances*, 11(6), 2021.

[15] Yuki Noguchi, Kei Matsushima, and Takayuki Yamada. Level set-based topology optimization for the design of labyrinthine acoustic metamaterials. *Materials & Design*, 219:110832, 2022.

[16] Peng Sheng, Xin Fang, Jihong Wen, and Dianlong Yu. Vibration properties and optimized design of a nonlinear acoustic metamaterial beam. *Journal of Sound and Vibration*, 492:115739, 2021.

[17] Hao-Wen Dong, Sheng-Dong Zhao, Peijun Wei, Li Cheng, Yue-Sheng Wang, and Chuanzeng Zhang. Systematic design and realization of double-negative acoustic metamaterials by topology optimization. *Acta Materialia*, 172:102–120, 2019.

[18] Steven A Cummer, Johan Christensen, and Andrea Alù. Controlling sound with acoustic metamaterials. *Nature Reviews Materials*, 1(3):1–13, 2016.

[19] Allan D Pierce. *Acoustics: an introduction to its physical principles and applications*. Springer, 2019.

[20] Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89):1–97, 2023.

[21] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.

[22] Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.

[23] Tapas Tripura and Souvik Chakraborty. Wavelet neural operator: a neural operator for parametric partial differential equations. *arXiv preprint arXiv:2205.02191*, 2022.

[24] Zhongkai Hao, Zhengyi Wang, Hang Su, Chengyang Ying, Yinpeng Dong, Songming Liu, Ze Cheng, Jian Song, and Jun Zhu. Gnot: A general neural operator transformer for operator learning. In *International Conference on Machine Learning*, pages 12556–12569. PMLR, 2023.

[25] Kamyar Azizzadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiaffini, Jean Kossaifi, and Anima Anandkumar. Neural operators for accelerating scientific simulations and design. *Nature Reviews Physics*, pages 1–9, 2024.

[26] Siddharth Nair, Timothy F Walsh, Greg Pickrell, and Fabio Semperlotti. Physics and geometry informed neural operator network with application to acoustic scattering. *arXiv preprint arXiv:2406.03407*, 2024.

[27] Steffen Marburg and Bodo Nolte. *Computational acoustics of noise propagation in fluids: finite and boundary element methods*, volume 578. Springer, 2008.

[28] COMSOL. Application gallery: Sonic crystal. https://www.comsol.com/model/sonic-crystal-16925. Accessed: 2024/08/28.

[29] Igor A. Baratta, Joseph P. Dean, Jørgen S. Dokken, Michal Habera, Jack S. Hale, Chris N. Richardson, Marie E. Rognes, Matthew W. Scroggs, Nathan Sime, and Garth N. Wells. Dolfinx: the next generation fenics problem solving environment. 2023.

[30] Ardavan F Oskooi, Lei Zhang, Yehuda Avniel, and Steven G Johnson. The failure of perfectly matched layers, and towards their redemption by adiabatic absorbers. *Optics Express*, 16(15):11376–11392, 2008.

[31] Christophe Geuzaine and Jean-François Remacle. Gmsh: A 3-d finite element mesh generator with built-in pre-and post-processing facilities. *International journal for numerical methods in engineering*, 79(11):1309–1331, 2009.

[32] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *The 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2623–2631, 2019.

[33] Joshua M. Long. Random fourier features pytorch. *GitHub*. *https://github.com/jmclong/random-fourier-features-pytorch*, 2021.

[34] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains, 2020.

A Implementation details

For the ground truth (FEM) solution, we simulate a domain consisting of 10 C-shaped structures aligned longitudinally in the x-direction, with a setup that matches the implementation in the COMSOL application gallery [28]. The domain is truncated in the x-direction using adiabatic layers [30] to ensure anechoic transmission for the crystal. It is truncated in the y-direction using hard sound boundaries, which is identical to periodic boundary conditions, due to the symmetry of the sonic crystal along the x-axis. We model the C-shapes with hard reflection properties, and choose the length of 10 unit cells in order to fully observe band gaps. All C-shapes have their opening pointing in the negative x-direction. We use a scattered field formulation [27] for the sound pressure to model a plane wave field incident normal to the crystal (traveling in x-direction).

For both the training and the test dataset, we uniformly sample the parameter space spanned by interval limits $R_O \leq 8.5\,\mathrm{mm}$, $R_I \leq 8.5\,\mathrm{mm}$, $b = 8.5\,\mathrm{mm}$, and a tolerance $\delta = 0.1\,\mathrm{mm}$. The unit cell has a constant width and height $d = 22\,\mathrm{mm}$. We use the tolerance to ensure that all sampled sonic crystals are valid, requiring $R_O \geq R_I + \delta$ and $R_I \geq b + \delta$. The parameter space is sampled 673 times for the training (90% train, 10% validation) and the test dataset 128 times.

We use Gmsh [31] to build meshes from the parameterizations. The meshes have at least 15 elements per wavelength at $20\,\mathrm{kHz}$. We compute the FEM solution for each observation on 256 equidistant frequencies $f \in [2\,\mathrm{kHz}, 20\,\mathrm{kHz}]$ using the DOLFINx computational environment of the FEniCSx framework [29]. Details for this procedure can be found on `Github`[3].

Our neural operators are implemented in our framework `continuiti`[4]. We aim to fix the numbers of all neural operators to approximately 10^5 parameters, and optimize the most important hyperparameter for our experiments, the batch size, using Optuna [32]. The internal structure of all our neural operators are fully connected NNs with residual connections, layer normalization, and hyperbolic tangent activation functions.

In order to allow the FNO to handle the parameter inputs, we use random Fourier features [33, 34] where we use Gaussian encoding with an embedding size 256 to encode the parameters of the sonic crystal. The FNO has a width of 8, depth 6, and we train it using a batch size of 4. The DON has a branch and trunk network with a width of 44, and a depth of 24 each. We use 64 basis functions, and a batch size of 4. The DNO has a width of 56, depth of 32 layers, and we train it using a batch size of 20. The DCO has a branch and trunk network with a with of 42, and a depth of 8 layers each, and the cat network is 32 layers deep with a width of 46. We train the DCO with a batch size of 20.

Each architecture is trained 10 times for the best parameters using different random seeds (but identical across different architectures). We employ Adam optimizer to minimize the mean squared error between prediction and ground truth values for 500 epochs. We reduce the learning rate from $\eta_0 = 1 \times 10^{-3}$ to $\eta_N = 1 \times 10^{-5}$ over the entire training process using cosine annealing. All samples undergo low-pass-filtering and normalization prior to being processed by the operators. The final operator is the best performing model with respect to the validation dataset. For greater detail, refer to `Neural-Operator-TL`[5] on Github. Training for 500 epochs on our machine takes on average $526\,\mathrm{s}$ for the DCO, $420\,\mathrm{s}$ for the DNO, $1130\,\mathrm{s}$ for the FNO, and $1934\,\mathrm{s}$ for the DON.

To evaluate the error metrics, we calculate the relative mean squared error across all models, using 10 different seeds. From this we employ bootstrapped statistics to get the 95% confidence intervals for the model performance.

In terms of hardware, we use an RTX 3060 Ti, an AMD Ryzen 7 2700x CPU and $16\,\mathrm{GB}$ of memory. To compare the speed of our approach to conventional methods, we use a COMSOL implementation [28], which takes $t = 21.6\,\mathrm{s}$ for 256 samples of the domain for one specific parameterization. The COMSOL implementation is faster than our FEniCSx code providing a reasonable state-of-the-art baseline, but in its current form it is not capable of solving many different parameterized geometries.

Figure 6 and Figure 7 show the transmission loss graph of two different configurations with both unfiltered and low-pass-filtered graphs. Figure 5 and Figure 8 display four frequency samples each for

[3] `https://github.com/JakobEliasWagner/Helmholtz-Sonic-Crystals`

[4] `https://github.com/aai-institute/continuiti`

[5] `https://github.com/JakobEliasWagner/Neural-Operator-TL`

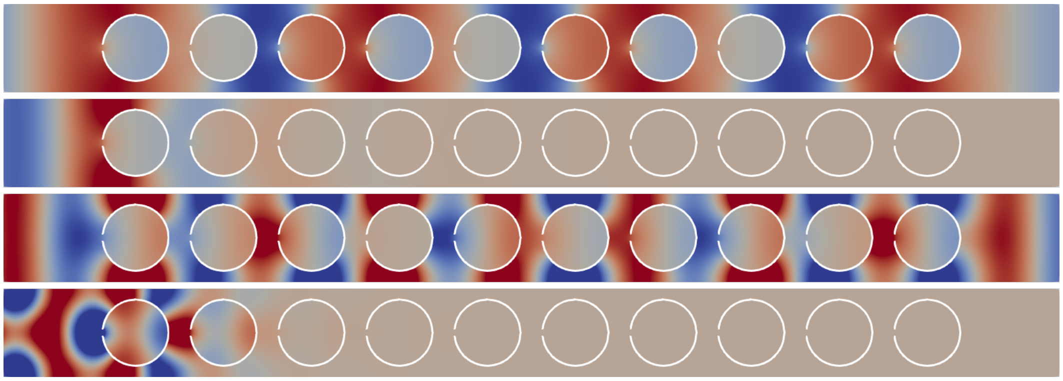

Figure 5: Real-valued pressure field for frequencies $f = \{4964\,\mathrm{Hz}, 7505\,\mathrm{Hz}, 11\,105\,\mathrm{Hz}, 17\,600\,\mathrm{Hz}\}$ (top to bottom) for the configuration $a = [8.5\,\mathrm{mm}, 8\,\mathrm{mm}, 1\,\mathrm{mm}]$.

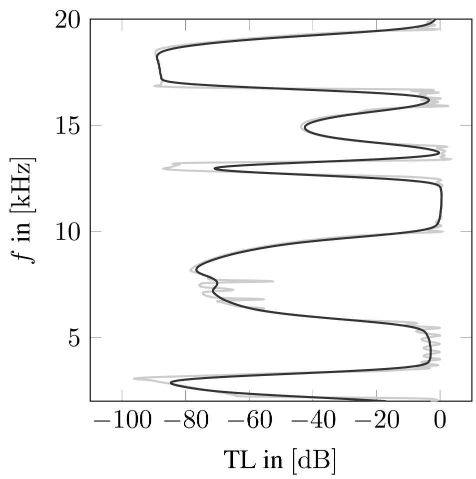

Figure 6: Unprocessed (light gray) and preprocessed (black, low-pass-filter) transmission loss graph for $a = [8.5\,\mathrm{mm}, 8\,\mathrm{mm}, 1\,\mathrm{mm}]$.

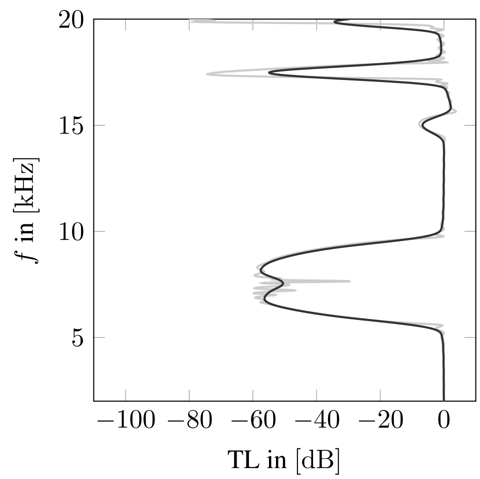

Figure 7: Unprocessed (light gray) and preprocessed (black, low-pass-filter) transmission loss graph for $a = [4\,\mathrm{mm}, 3.7\,\mathrm{mm}, 3\,\mathrm{mm}]$.

both configurations. The band gaps in C-shaped sonic crystals stem from two different phenomena: Bragg-scattering (crystalline structure) and local resonance of the inner chamber of the C-shapes [9].

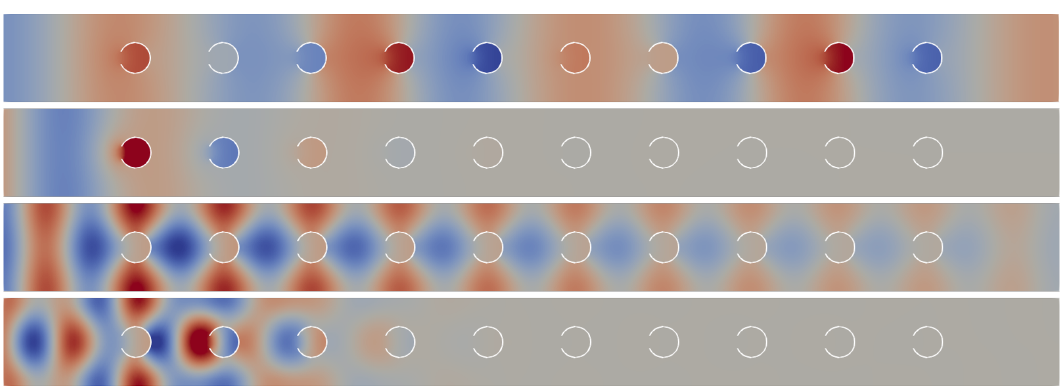

Figure 8: Real-valued pressure field for frequencies $f = \{5105\,\mathrm{Hz}, 7505\,\mathrm{Hz}, 15\,058\,\mathrm{Hz}, 20\,000\,\mathrm{Hz}\}$ (top to bottom) for the configuration $a = [4\,\mathrm{mm}, 3.7\,\mathrm{mm}, 3\,\mathrm{mm}]$.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The problem setting and methodology are clearly outlined with technical details, and the paper provides extensive empirical data that supports the claims made in the abstract and introduction.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper acknowledges its limitations, such as the focus on a single dataset and problem, and the absence of comparisons to classical surrogate models. These limitations are clearly addressed and suggest directions for future research.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory Assumptions and Proofs**

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Assumptions for the new architectures and the experiment can be found in detail and the entire process, including the code, is available.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The experimental results can be recreated using the provided code and is described in detail in the appendix of the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: Access to our framework, data generation, and the specific implementations for this paper are available.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental Setting/Details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: Dataset size, data splits, hyperparameters, the type of optimizer and the training process itself is described in the paper.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment Statistical Significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: Both for plots and in the table confidence intervals are available.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments Compute Resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: The details of the experiment include information on the compute resources and training and inference time.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code Of Ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: The research conforms with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [NA]

    Justification: No societal impact of the work was performed.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: The paper poses no such risk.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: The code, data and models are labeled properly.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: We use artificially created data using an open source framework (FEniCSx).

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.