

A learned Born series for highly-scattering media

Antonio Stanziola, Simon Arridge,  Ben T. Cox,  and Bradley E. Treeby 

University College London, Gower Street, London, WC1E 6BT, United Kingdom

a.stanziola@ucl.ac.uk, s.arridge@ucl.ac.uk, b.cox@ucl.ac.uk, b.treeby@ucl.ac.uk

Abstract: A new method for solving the wave equation is presented, called the learned Born series (LBS), which is derived from a convergent Born series but its components are found through training. The LBS is shown to be significantly more accurate than the convergent Born series for the same number of iterations, in the presence of high contrast scatterers, while maintaining a comparable computational complexity. The LBS is able to generate a reasonable prediction of the global pressure field with a small number of iterations, and the errors decrease with the number of learned iterations. © 2023 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).

[Editor: D. Keith Wilson]

<https://doi.org/10.1121/10.0017937>

Received: 13 February 2023 **Accepted:** 5 April 2023 **Published Online:** 1 May 2023

1. Introduction

Computational solvers for the Helmholtz equation with heterogeneous material properties are used in many fields that require single frequency wave simulations, including optics, seismology, acoustics, and electromagnetics. One approach is to compute the terms of the Born series (Born, 1926) through the recursive application of the known Green's operator for a homogeneous medium. However, in its conventional form, the Born series converges only for low contrast scattering potentials—a major limitation in many practical applications. As an example, the application that has motivated this work is the propagation of ultrasound waves through the skull, which has a sound speed about double that of the surrounding soft tissue. A convergent Born series (CBS) was proposed by Kleinman *et al.* (1988, 1990) based on a generalised overrelaxation method. More recently, Osnabrugge *et al.* (2016) introduced another CBS through the use of a preconditioner based on the scattering potential, which has since found many applications (Jakobsen *et al.*, 2020; Kaushik *et al.*, 2020; Krüger *et al.*, 2017; Lee *et al.*, 2022). However, for strongly scattering media many iterations are needed, as the convergence rate is limited by the range of spatial wavenumbers in the problem. To overcome the slow convergence rate, we present a learned Born series (LBS), in which a generalised preconditioner and the modified Green's operator are found, for a given number of unrolled iterations (Monga *et al.*, 2021), through training. We show that the convergence rate in the presence of high contrast scatterers is considerably improved.

2. Learned Born series

We aim to solve the Helmholtz differential equation $(\nabla^2 + k^2)u = -s$, where $k = \omega/c_0$ is the heterogeneous wavenumber, ω is the angular frequency, and c_0 is the spatially-varying sound speed. To solve this equation, following Osnabrugge *et al.* (2016), we start by defining the scattering potential $v = k^2 - \kappa^2 - i\varepsilon$, where $\kappa^2, \varepsilon \in \mathbb{R}_+$ are positive scalars, and rewrite the differential equation as

$$\mathcal{L}u = (\nabla^2 + \kappa^2 + i\varepsilon)u = -s - vu. \quad (1)$$

We denote with \mathcal{G} the Green's operator associated with the operator \mathcal{L} , that is,

$$\mathcal{G} = \mathcal{F}^{-1}(p^2 - \kappa^2 - i\varepsilon)^{-1}\mathcal{F}, \quad (2)$$

where \mathcal{F} is the Fourier transform operator and $p^2 = p_x^2 + p_y^2$ is the squared magnitude of the Fourier transformed coordinates. Applying the Green's operator to both sides of Eq. (1) leads to $u = \mathcal{G}s + \mathcal{G}vu$. Repeated substitution of the right-hand-side into u leads to the classical form of the Born series,

$$u = [1 + \mathcal{G}v + (\mathcal{G}v)^2 + (\mathcal{G}v)^3 + \dots]\mathcal{G}s = \sum_{n=0}^{\infty} (\mathcal{G}v)^n \mathcal{G}s, \quad (3)$$

under the assumption that the series converges. As mentioned in Sec. 1, this assumption is satisfied only for low scattering media. Osnabrugge *et al.* (2016) proposed the preconditioner $q = iv/\varepsilon$, leading to the CBS

$$u = [1 + \mathcal{M} + \mathcal{M}^2 + \mathcal{M}^3 + \dots]q\mathcal{G}s, \quad \mathcal{M} = q\mathcal{G}v + (1 - q), \quad (4)$$

which is convergent for scattering media with arbitrarily-strong sound speed contrasts for an appropriately chosen ε . Despite the convergence guarantees of Eq. (4), the convergence rate is inversely proportional to the maximum wavenumber difference between any two points in the domain. This means that for media with sound-speed high contrast, a large number of iterations are necessary.

With the aim of reducing the number of iterations required to solve problems in highly-scattering media, we modify the CBS as follows. First, note that the series in Eq. (4) can be rewritten as

$$u_n = u_{n-1} + w_n, \quad w_n = \mathcal{M}w_{n-1}, \quad (5)$$

with $u_0 = w_0 = q\mathcal{G}s$. To develop a learned version of Eq. (4), we transform the linear operator \mathcal{M} into a non-linear one in the form

$$\mathcal{M} = q\mathcal{G}v + (1 - q) \rightarrow \widehat{\mathcal{M}}_\theta = (A_\theta \widehat{\mathcal{G}}_\theta B_\theta + C_\theta). \quad (6)$$

Here, A is the learned version of the preconditioner q , B is the learned version of the scattering potential v , and C is now a general additive term (that is not enforced to be $1 - A$). Each of A , B , and C are the outputs of a neural network dependent on the input sound speed and the spatial location maps, i.e.,

$$A_\theta = f^A(c, \mathbf{x}; \theta), \quad B_\theta = f^B(c, \mathbf{x}; \theta), \quad C_\theta = f^C(c, \mathbf{x}; \theta), \quad (7)$$

where, to simplify notation, we use θ to denote the learnable parameters for each operator at each iteration, without implying that they are the same set. The learned fields (A, B, C) are returned by three fully connected neural networks (FCNN) acting channel wise on the (c, \mathbf{x}) map, also known as 1×1 convolutions, with two hidden layers. The size of the hidden layers is equal to the number of channels. Similar to convolutional networks, we generalise Eq. (6) to c channels by associating the action of each (A, B, C) field to a matrix product; for example, if $\mathbf{u} \in \mathbb{R}^c = (u_1, u_2, \dots, u_c)$, then $(A\mathbf{u})_i = \sum_j A_{ij}u_j$. Note that the values of (A, B, C) change for each pixel location.

We also make the wavenumber shift of the Green's function, $\kappa^2 + i\varepsilon$, learnable. To avoid dividing by zero during training, instead of directly learning κ^2 and ε , we parameterise the Green's operator as

$$\widehat{\mathcal{G}}_\theta = \mathcal{F}^{-1}(\alpha p^2 - \kappa^2 - i)^{-1} \mathcal{F}, \quad \theta = \{\alpha, \kappa^2\}, \quad (8)$$

where α and κ^2 are learnable parameters, which is equivalent up to a scaling factor (that is absorbed by the A term). Similarly, the input to the Green's operator is an n -dimensional field, therefore the Green's function is a mapping $\mathbb{R}^c \rightarrow \mathbb{R}^c$, thus α, κ^2 are $c \times c$, where c is the number of channels. Last, we apply a projected non-linearity of the form $\widehat{\sigma}(u) = D_1\sigma(D_2u)$, where D_1, D_2 are two dense layers acting channel-wise and σ is a non-linear function, in our case a GeLU function (Hendrycks and Gimpel, 2016). The final learned iteration is therefore

$$u_n = u_{n-1} + w_n, \quad w_n = \widehat{\sigma}_n(\widehat{\mathcal{M}}_n w_{n-1}), \quad n = 1, \dots, m, \quad (9)$$

where we have used the subscript n to highlight that $\widehat{\sigma}$ and $\widehat{\mathcal{M}}$ have different parameters at each iteration. We set $u_0 = 0$. We note the similarity between Eq. (9) to the Fourier Neural Operators (FNO) framework (Li et al., 2020), where the network is set up as the sum of a Fourier filter and a local convolution operator as

$$u_{n+1} = \sigma(\mathcal{N}u_n), \quad \mathcal{N} = \mathcal{F}^{-1}K\mathcal{F} + C, \quad (10)$$

where K is a learnable Fourier filter and C is a learnable local convolution filter. Similar to the FNO, we pad the input by a fixed amount to accommodate the artifacts generated by the periodicity assumed by the discrete Fourier transform and afterwards unpad the output. A block diagram of the network is given in Fig. 1.

3. Training and results

A dataset was constructed from 2000 MNIST (Deng, 2012) samples upsampled to 128×128 pixels, which were rescaled and used as sound speed distributions. MNIST is a dataset of binary images of the digits 0 to 9 which is widely used in machine learning research. Here, we use it to create a diverse range of highly scattering spatial distributions. The two-




Fig. 1. Block diagram for the learned Born series (LBS) with n iterations.

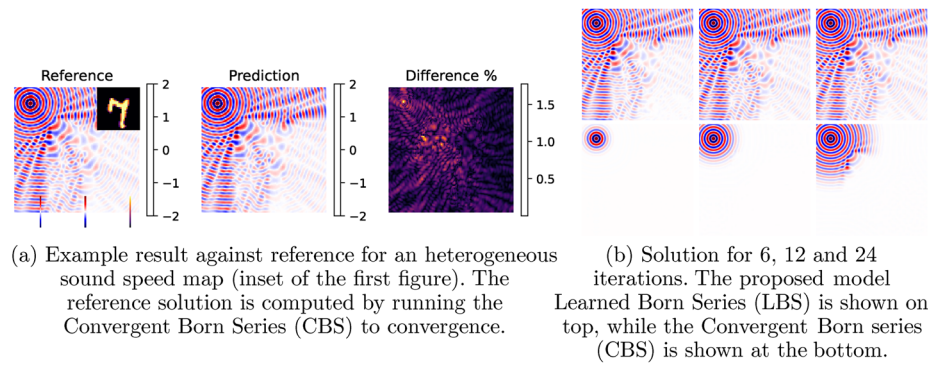


Fig. 2. Example results on the test set.

dimensional (2D) simulations were performed in normalised units with a background sound speed of 1 m.s^{-1} , maximum sound speed of 2 m.s^{-1} , and $\omega = 1 \text{ rad.s}^{-1}$ (giving 2π grid points per wavelength). The dataset was split between train, validation, and test in 80%, 10%, and 10%-sized subsets. For each sound speed map, a reference solution was generated using the Helmholtz solver in j-Wave (Stanziola *et al.*, 2023) using GMRES with a fixed source position. The loss function was the mean squared error against the reference solution. All networks were trained with a batch-size of 16 and learning rate of $3 \cdot 10^{-3}$: the network corresponding to the epoch with the lowest validation loss was saved, although we did not observe any over-fitting during training. The network was trained for 6, 12, and 24 unrolled iterations, with $c=8$. Here, by unrolled iterations, we mean the value of n in Fig. 1, which is analogous to the number of leading terms retained in the truncated scattering series. We also trained an extremely light version with $c=2$ (matching the number of components on a complex field, as in the original CBS) with linear activation functions.

The simulated pressure field for a representative example from the test set using a network trained with 12 iterations is shown in Fig. 2(a). Even with a very small number of iterations, the LBS is able to generate a reasonable prediction of the global pressure field. The errors against j-Wave for the whole test set are shown in Fig. 3(a). As expected, the errors decrease with the number of learned iterations. Compared to the CBS run for the same number of iterations, the error is significantly reduced. This is because the pseudo propagation speed in the CBS is limited due to the strong contrast, as shown in Fig. 3(b).

The outliers correspond to the speed of sound maps that exhibit resonant modes, e.g., “0” digits for which the field exhibited a mode-like structure similar to whispering gallery modes (Stanziola *et al.*, 2021). The error for the CBS was also proportionally higher for these examples. To match the same error on the test set of the LBS using 6, 12, and 24 iterations, on average the CBS required 13.1, 8.6, and 5.0 times more iterations, respectively [see Fig. 3(b)]. Even when the number of channels is reduced to 2 (equivalent to the complexity of the CBS), the convergence of the LBS is still much faster than the CBS [see yellow boxplot in Fig. 3(a)].

4. Summary and discussion

We present a learned Born series, inspired by the convergent Born series presented in Osnabrugge *et al.* (2016), that significantly improves convergence rates for single-frequency wave simulations with strong sound speed contrast. While for channel counts $c > 2$ the LBS requires more fast Fourier transforms (FFTs) at each layer compared to the CBS, such FFTs

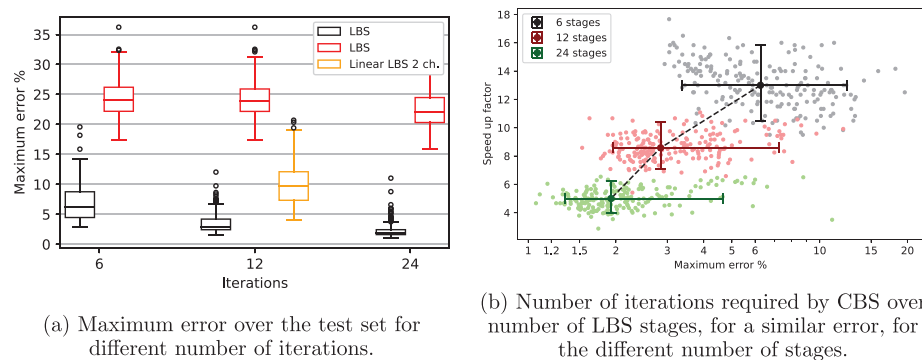


Fig. 3. Error and speed comparison against the CBS on the test set.

can be computed in a parallel fashion, which still makes the method more efficient in practice. While we have only shown results for a single source position, results for more complex sources can be obtained by superposition, or by retraining with a given source geometry; see Stanzola *et al.*, 2021. The code for reproducing our results is available online in Stanzola, 2023.

In the future, it would be interesting to explore several potential extensions of this work. One possibility is to train on sub-sampled problems that are otherwise unsolvable with classical methods in order to further reduce the computational complexity of the LBS. This approach could have applications in areas such as transcranial treatment planning and full waveform inversion. Another interesting direction would be to apply the method to other partial differential equations (PDEs) for which a scattering potential can be defined, such as the Schrödinger equation.

One potential way to improve the LBS would be to substitute the fixed padding with a learnable “perfectly matched layer” layer, instead of feeding the spatial grid \mathbf{x} . This could lead to an architecture that is resolution and grid-size independent, making it a neural operator as described in Kovachki *et al.* (2021). It would also be interesting to train the method as an iterative solver, as described in Stanzola *et al.* (2021). Finally, it would be interesting to explore whether the LBS can be extended to heterogeneous density. While the method presented in this paper does not support it, the CBS does not either. Allowing so would open up new applications in areas where the density of the medium being modeled is nonuniform, such as ultrasound applications with the presence of bone or geophysics.

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC), UK, Grant Nos. EP/V026259/1, EP/S026371/1, EP/T022280/1, EP/W029324/1, and EP/T014369/1.

References and links

- Born, M. (1926). “Quantenmechanik der stoßvorgänge” (“Quantum mechanics of collision processes”), *Z Phys.* **38**(11), 803–827.
- Deng, L. (2012). “The MNIST database of handwritten digit images for machine learning research,” *IEEE Signal Process. Mag.* **29**(6), 141–142.
- Hendrycks, D., and Gimpel, K. (2016). “Gaussian error linear units (GELUS),” [arXiv:1606.08415](https://arxiv.org/abs/1606.08415).
- Jakobsen, M., Huang, X., and Wu, R.-S. (2020). “Homotopy analysis of the Lippmann–Schwinger equation for seismic wavefield modeling in strongly scattering media,” *Geophys. J. Int.* **222**(2), 743–753.
- Kaushik, A., Yalavarthy, P. K., and Saha, R. K. (2020). “Convergent Born series improves the accuracy of numerical solution of time-independent photoacoustic wave equation,” *J. Mod. Opt.* **67**(9), 849–855.
- Kleinman, R., Roach, G., Schuetz, L., and Shirron, J. (1988). “An iterative solution to acoustic scattering by rigid objects,” *J. Acoust. Soc. Am.* **84**(1), 385–391.
- Kleinman, R., Roach, G., and Van Den Berg, P. (1990). “Convergent Born series for large refractive indices,” *J. Opt. Soc. Am. A* **7**(5), 890–897.
- Kovachki, N., Li, Z., Liu, B., Azizzadenesheli, K., Bhattacharya, K., Stuart, A., and Anandkumar, A. (2021). “Neural operator: Learning maps between function spaces,” [arXiv:2108.08481](https://arxiv.org/abs/2108.08481).
- Krüger, B., Brenner, T., and Kienle, A. (2017). “Solution of the inhomogeneous Maxwell’s equations using a Born series,” *Opt. Express* **25**(21), 25165–25182.
- Lee, M., Hugonnet, H., and Park, Y. (2022). “Inverse problem solver for multiple light scattering using modified Born series,” *Optica* **9**(2), 177–182.
- Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., and Anandkumar, A. (2020). “Fourier neural operator for parametric partial differential equations,” [arXiv:2010.08895](https://arxiv.org/abs/2010.08895).
- Monga, V., Li, Y., and Eldar, Y. C. (2021). “Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing,” *IEEE Signal Process. Mag.* **38**(2), 18–44.
- Osnabrugge, G., Leedumrongwatthanakun, S., and Vellekoop, I. M. (2016). “A convergent Born series for solving the inhomogeneous Helmholtz equation in arbitrarily large media,” *J. Comput. Phys.* **322**, 113–124.
- Stanzola, A. (2023). “Learned Born series (v1.0) [Code],” <https://github.com/ucl-bug/lbs> (Last viewed: 5 April 2023).
- Stanzola, A., Arridge, S. R., Cox, B. T., and Treeby, B. E. (2021). “A Helmholtz equation solver using unsupervised learning: Application to transcranial ultrasound,” *J. Comput. Phys.* **441**, 110430.
- Stanzola, A., Arridge, S. R., Cox, B. T., and Treeby, B. E. (2023). “j-Wave: An open-source differentiable wave simulator,” *SoftwareX* **22**, 101338.