Calculating coherent light-wave propagation in large heterogeneous media

Vettenburg, Tom; Horsley, Simon A. R.; Bertolotti, J.

Publication date:
2018

Document Version
Early version, also known as pre-print

Calculating coherent light-wave propagation in large heterogeneous media

T. Vettenburg[1,2,*], S.A.R. Horsley[1], and J. Bertolotti[1]

[1] *Department of Physics and Astronomy, University of Exeter, Exeter, EX4 4QL, United Kingdom*
[2] *School of Science and Engineering, University of Dundee, Nethergate, Dundee, DD1 4HN, United Kingdom*
[*] *Corresponding author: t.vettenburg@dundee.ac.uk*

Abstract

Understanding the interaction of light with a highly scattering material is essential for optical microscopy of optically thick and heterogeneous biological tissues. Ensemble-averaged analytic solutions cannot provide more than general predictions for relatively simple cases. Yet, biological tissues contain chiral organic molecules and many of the cells' structures are birefringent, a property exploited by polarization microscopy for label-free imaging. Solving Maxwell's equations in such materials is a notoriously hard problem. Here we present an efficient method to determine the propagation of electro-magnetic waves in arbitrary anisotropic materials. We demonstrate how the algorithm enables large scale calculations of the scattered light field in complex birefringent materials, chiral media, and even materials with a negative refractive index.

1 Introduction

Determining how light propagates in heterogeneous media is a notoriously hard problem[1]. Unless the system of interest has symmetries such as periodicity[2], one needs to solve the Maxwell equations *ab initio* with appropriate boundary conditions. While this may be feasible for relatively simple systems such as Mie scattering[3], multiple scattering of light can lead to many subtle effects[4–6]. The fractal propagation method can simulate light in biological tissues[7]; however, calculating the exact light field distribution in arbitrary large heterogeneous materials remains out of reach for the current generation of computational methods[1]. A further complication is that the media are generally not isotropic, meaning that the refractive index is different depending on the orientation of the field polarization. Such birefringence is common in many samples of interest such as TiO_2, the lipid bilayer cell-membrane[8–10], or muscle fibers[11]. Indeed, birefringence is a valuable, label-free, contrast method[11–16]. Furthermore, biological tissues contain chiral organic molecules such as glucose, a property that is linked to the electro-magnetic coupling and cannot be modeled by anisotropic permittivity. In this paper we investigate the application of the modified Born series[17,18] to solve Maxwell's equations in large heterogeneous electromagnetic media, characterized by arbitrary linear constitutive relations.

Although finite-element and finite-difference time-domain (FDTD) methods can in principle handle general electromagnetic problems of an arbitrary size, such methods do not scale well to the dimensions relevant in microscopy. Only recently it has become possible to represent the complete field distribution in computer memory for larger samples. Conventional methods require amounts of high-speed access storage that are considerably larger. For P sample points the finite-element method typically requires the representation of a sparse $3P \times 3P$ matrix in memory[19], which is then iteratively applied in order to reach the desired solution. Meanwhile, the FDTD method requires far higher sampling densities to limit the accumulation of errors. A quite different, and more physically meaningful iterative method is the Born series, which calculates the scattered field, with each iteration including the effect of the next highest order of scattering events. This approach is commonly used for the analytical theory of multiple scattering[20,21]. When implemented numerically, all operations can be performed in a constant space ($\propto P$) and using only fast-Fourier transforms and operations on diagonal matrices. The high computational efficiency and modest memory requirements make the Born series an ideal candidate for solving large-scale electromagnetic problems. However, a severe limitation is that the basic form of the series only converges in the limit of weak scattering. A modified Born series for scalar waves, which converges for any value of the scattering strength, was recently proposed by Osnabrugge et al.[17], and generalized to vector waves by Krüger et al.[18]. Yet, as it stands, this approach is limited to media with an isotropic permittivity and without magnetic properties. The modified Born series method thus excludes a large class of materials such as biological tissues that exhibit birefringence or contain chiral organic compounds, as well as any potential applications to metamaterial structures. Here we generalize the numerical Born series method to arbitrary linear materials, including those with heterogeneous magnetic properties and bi-anisotropy. The modification to the Born series is found to be notably more subtle in these materials. We demonstrate that the algorithm enables large scale calculations of the scattered light field in complex birefringent materials, chiral media, and even materials with negative refractive index. Furthermore, we show that the iteration is robust to numerical errors.

To start we review the simplest application of the Born series, and its modified form, as given in[17,18]. At a fixed frequency ω, the electric and magnetic fields satisfy Maxwell's equations, $\boldsymbol{\nabla} \times \boldsymbol{E}(\boldsymbol{x}) = \mathrm{i}\omega \boldsymbol{B}(\boldsymbol{x})$, and $\boldsymbol{\nabla} \times \boldsymbol{H}(\boldsymbol{x}) = \boldsymbol{j}(\boldsymbol{x}) - \mathrm{i}\omega \boldsymbol{D}(\boldsymbol{x})$, where $\boldsymbol{j}(\boldsymbol{x})$ is the electric current density as a function of the spatial coordinate, $\boldsymbol{x}$. In first instance, we consider the case that the medium is isotropic and non-magnetic, and can thus be characterized in terms of a scalar relative permittivity $\epsilon(\boldsymbol{x})$. This quantity connects the displacement field $\boldsymbol{D}$ and magnetic flux density $\boldsymbol{B}$ fields to the electric field $\boldsymbol{E}$ and magnetic field $\boldsymbol{H}$ in the constitutive relations:

$$\boldsymbol{D}(\boldsymbol{x}) = \epsilon_0 \epsilon(\boldsymbol{x}) \boldsymbol{E}(\boldsymbol{x}), \quad \text{and} \quad \boldsymbol{B}(\boldsymbol{x}) = \mu_0 \boldsymbol{H}(\boldsymbol{x}) \tag{1}$$

where ϵ_0 and μ_0 are the vacuum permittivity and permeability, respectively. By substituting both constitutive relations into Maxwell's equations we obtain the vector Helmholtz equation for the electric field

$$\boldsymbol{\nabla} \times \boldsymbol{\nabla} \times \boldsymbol{E}(\boldsymbol{x}) - k_0^2 \epsilon(\boldsymbol{x}) \boldsymbol{E}(\boldsymbol{x}) = \mathrm{i}\omega\mu_0 \boldsymbol{j}(\boldsymbol{x}) = \boldsymbol{S}(\boldsymbol{x}), \tag{2}$$

where $k_0 = \omega/c$ is the free space wavenumber and $\boldsymbol{S}(\boldsymbol{x})$ represents the radiation source. In order to solve equation (2)

for the electric field $\boldsymbol{E}(\boldsymbol{x})$, one must invert the operator $\boldsymbol{\mathcal{O}} = -k_0^2\epsilon(\boldsymbol{x}) + \boldsymbol{\nabla}\times\boldsymbol{\nabla}\times$. We could attempt to do this directly, representing $\boldsymbol{\mathcal{O}}$ as a matrix in some basis of eigenfunctions and then applying a numerical inversion algorithm. However, for moderately large systems the direct representation and inversion of such matrix is infeasible, due to demands of both memory and time. Iterative solutions are thus required where, starting from some initial guess for the electric field, one repeatedly applies a matrix until the result is arbitrarily close to the one that would be obtained from a direct inversion, as for instance in finite element simulations (see e.g.[19], Chapter 19).

The Born series[22] is a physically motivated version of this iterative procedure that was used as a mathematical technique long before the current numerical approaches[20]. It involves splitting the Helmholtz operator, $\boldsymbol{\mathcal{O}}$, into two parts $\boldsymbol{\mathcal{O}} = \boldsymbol{\mathcal{O}}_h + \boldsymbol{\mathcal{O}}_i$, where the inverse of $\boldsymbol{\mathcal{O}}_h$ is known, expanding the rest as what is known as the Born series:

$$\begin{aligned}\boldsymbol{\mathcal{O}}^{-1} = (\boldsymbol{\mathcal{O}}_h + \boldsymbol{\mathcal{O}}_i)^{-1} &= \left(\mathbb{1}_3 + \boldsymbol{\mathcal{O}}_h^{-1}\boldsymbol{\mathcal{O}}_i\right)^{-1}\boldsymbol{\mathcal{O}}_h^{-1} \\ &= \left[\mathbb{1}_3 - \left(\boldsymbol{\mathcal{O}}_h^{-1}\boldsymbol{\mathcal{O}}_i\right) + \left(\boldsymbol{\mathcal{O}}_h^{-1}\boldsymbol{\mathcal{O}}_i\right)^2 + \dots\right]\boldsymbol{\mathcal{O}}_h^{-1} \\ &= \left[\sum_{p=0}^{\infty}\left(-\boldsymbol{\mathcal{O}}_h^{-1}\boldsymbol{\mathcal{O}}_i\right)^p\right]\boldsymbol{\mathcal{O}}_h^{-1}, \quad (3)\end{aligned}$$

where $\mathbb{1}_3$ is the identity operator in three-dimensional space. The operator $\boldsymbol{\mathcal{O}}_h$ is typically associated with propagation through a homogeneous medium, and $\boldsymbol{\mathcal{O}}_i$ is due to the inhomogeneity of the material. We can thus understand the Born series (3) as representing a sequence of scattering events, where as our initial guess the source $\boldsymbol{S}$ generates the same field as if the material were homogeneous $\boldsymbol{E}_0(\boldsymbol{x}) = \boldsymbol{\mathcal{O}}_h^{-1}\boldsymbol{S}$, the first iteration adds to this events where a single scattering event takes place (proportional to $\boldsymbol{\mathcal{O}}_h^{-1}\boldsymbol{\mathcal{O}}_i$), the second iteration introduces double scattering (proportional to $(\boldsymbol{\mathcal{O}}_h^{-1}\boldsymbol{\mathcal{O}}_i)^2$), and so on. The main problem with this intuitive expansion is that in many cases it will not converge. Physically, this is connected to the existence of bound states[22], which are due to the constructive interference of an infinite number of scattering events. On the other hand, the mathematical origin of this divergence is simple, it is the same as the divergence of the scalar series $(1-q)^{-1} = \sum_{p=0}^{\infty} q^p, \forall q : q \in \mathbb{C}$: for convergence we must require that $|q| < 1$. Correspondingly, the Born series, as given by (3), will only converge if all of the eigenvalues of $\boldsymbol{\mathcal{O}}_h^{-1}\boldsymbol{\mathcal{O}}_i$ have a magnitude less than unity (for discussion see[22]). Although there are cases where exact results can be deduced from this series[23], the application of the Born series is generally restricted to weakly scattering media.

For our particular case we take the homogeneous medium of $\boldsymbol{\mathcal{O}}_h$ to have a constant permittivity $\alpha = \alpha_\mathrm{r} + \mathrm{i}\alpha_\mathrm{i}$. The inverse of this operator is given by the Green function for this medium, $\boldsymbol{\mathcal{O}}_h^{-1} = \overleftrightarrow{G}$, which is the two-index object (dyadic) that is the solution to

$$\boldsymbol{\nabla}\times\boldsymbol{\nabla}\times\overleftrightarrow{\boldsymbol{G}}(\boldsymbol{x},\boldsymbol{x}') - k_0^2\alpha\overleftrightarrow{\boldsymbol{G}}(\boldsymbol{x},\boldsymbol{x}') = \mathbb{1}_3\delta^{(3)}(\boldsymbol{x}-\boldsymbol{x}'). \quad (4)$$

The remaining part of our operator $\boldsymbol{\mathcal{O}}$ (2), is the spatially-variant difference $\boldsymbol{\mathcal{O}}_i = -k_0^2(\epsilon(\boldsymbol{x}) - \alpha) \equiv -k_0^2\chi(\boldsymbol{x})$, where $\chi(\boldsymbol{x})$ is the isotropic susceptibility with respect to our choice of background permittivity, $\alpha \in \mathbb{C}$. The Born series, as defined in (3), for the electric field $\boldsymbol{E}$ then becomes

$$\boldsymbol{E} = \boldsymbol{\mathcal{O}}^{-1}\boldsymbol{S} = \left[\sum_{p=0}^{\infty}(k_0^2\overleftrightarrow{\boldsymbol{G}}\chi)^p\right]\overleftrightarrow{\boldsymbol{G}}\boldsymbol{S}. \quad (5)$$

Although we have not yet addressed the problem of convergence, the numerical advantage of this iterative series is that it can be done in a constant space and with a combination of fast-Fourier Transform algorithm[24]) and operations on matrices that are diagonal in their spatial indices. To see this, we observe that χ is diagonal in its spatial indices, and that the dyadic Green function (the solution to equation (4)) can be written as a product of a Fourier transform $\mathcal{F}$, a diagonal matrix (in Fourier space) and an inverse Fourier transform (see supplementary section S1B)

$$\overleftrightarrow{\boldsymbol{G}} = \mathcal{F}^{-1}\left(\frac{\boldsymbol{\Pi}_\mathrm{T}}{\boldsymbol{k}^2 - \alpha k_0^2} - \frac{\boldsymbol{\Pi}_\mathrm{L}}{k_0^2\alpha}\right)\mathcal{F} \quad (6)$$

where $\boldsymbol{\Pi}_\mathrm{T}$ and $\boldsymbol{\Pi}_\mathrm{L}$ are projection matrices that separate the field in a super-position of plane waves with transverse and longitudinal electric-components, respectively. The projection operations are defined as the outer product of the normalized k-vectors, $\boldsymbol{\Pi}_\mathrm{L} = \boldsymbol{k}\otimes\boldsymbol{k}/\|\boldsymbol{k}\|^2$, and $\boldsymbol{\Pi}_\mathrm{T} = \mathbb{1}_3 - \boldsymbol{\Pi}_\mathrm{L}$. The quantity in the rounded brackets can thus be seen to be a diagonal matrix in the two Fourier space indices.

As mentioned above, in a pair of recent papers[17,18], it was shown that this Born series (5) can be made into a convergent series, the only proviso being that the system exhibits solely dissipation (no gain), and that the permittivity is nowhere singular. To achieve this convergence the authors used two ingenious steps. The first is to notice that the complex background permittivity α can be chosen freely, changing the eigenvalues of the operator $k_0^2\overleftrightarrow{\boldsymbol{G}}\chi$, without affecting the solution. The second is that one can apply a pre-conditioner, Γ, to both sides of the recursion relation (5) as follows:

$$\begin{aligned}\Gamma\boldsymbol{E} &= k_0^2\Gamma\overleftrightarrow{\boldsymbol{G}}\chi\boldsymbol{E} + \Gamma\overleftrightarrow{\boldsymbol{G}}\boldsymbol{S} \\ \rightarrow\boldsymbol{E} &= (k_0^2\Gamma\overleftrightarrow{\boldsymbol{G}}\chi + \mathbb{1}_3 - \Gamma)\boldsymbol{E} + \Gamma\overleftrightarrow{\boldsymbol{G}}\boldsymbol{S} \quad (7)\end{aligned}$$

The corresponding so-called 'modified Born series', which is the counterpart of series (5), is then given by

$$\boldsymbol{E} = \left[\sum_{p=0}^{\infty}\boldsymbol{M}^p\right]\Gamma\overleftrightarrow{\boldsymbol{G}}\boldsymbol{S}, \quad \text{where } \boldsymbol{M} \equiv k_0^2\Gamma\overleftrightarrow{\boldsymbol{G}}\chi - \Gamma + \mathbb{1}_3, \quad (8)$$

which converges when the absolute value of the largest eigenvalue of $\boldsymbol{M}$ is less than 1. This condition has been shown[17,18] to hold when both the preconditioning operator is given as $\Gamma = \frac{\mathrm{i}}{\alpha_\mathrm{i}}\chi$ and when the imaginary part of the background permittivity, α_i, is larger than the largest value of $|\Delta(\boldsymbol{x})| = |\epsilon(\boldsymbol{x}) - \alpha_\mathrm{r}|$, considered over all points, $\boldsymbol{x}$, in space. The aforementioned work showed that this modified series can be used to very quickly compute the electromagnetic field in large media of moderate index contrast. This result relies on the isotropy of the permittivity, $\epsilon(\boldsymbol{x})$ and on the absence of any magnetic effect or chirality. In this paper we generalize this approach to media with any permittivity or permeability, including anisotropic, magnetic, chiral, and bi-anisotropic media. Interestingly we find that proving the convergence of the modified series seems to be much more subtle when these additional optical effects are accounted for. We also discuss the convergence speed and memory requirements of this approach, and present an easy to implement algorithm to simulate electromagnetic wave propagation in an arbitrary medium. The method is implemented as a *Python 3* library and available with examples for download at `https://uod.box.com/v/macromax`.

The paper is organized as follows. In section 2 we describe how the modified Born series method can be generalized to anisotropic dielectrics. The iterative algorithm is introduced and we discuss how the preconditioner values can be calculated to ensure convergence. The algorithm is demonstrated with the polarization-dependent propagation through

homogeneous and highly heterogeneous birefringent materials. Next, we show that the method is not limited to Hermitian permittivity matrix-functions and that it can be extended to non-Hermitian permittivity. In Section 3 we show how the preconditioning step of the algorithm can be adapted so that the same iteration can be used for magnetic materials. Accounting for electric-magnetic coupling also enables chiral, Tellegen[25], and general bi-anisotropic media. This is demonstrated with the propagation of a linearly polarized wave of visible light through 10 mm of a highly chiral substance. Finally we show how the method naturally handles more esoteric materials such as a solid with a negative refractive index, thus demonstrating the capability of simulating light propagation in such metamaterials.

2 A convergent Born series for anisotropic dielectrics

As a first step towards a Born series for arbitrary electromagnetic materials, we consider non-magnetic birefringent materials. Optical elements such as waveplates and Wollaston prisms are birefringent and characterized by anisotropic permittivity that cannot be represented by a scalar function, $\epsilon(\boldsymbol{x})$. In this section we show how the modified Born series can be extended to permittivity tensors. Later it will be shown that arbitrary electromagnetic problems can be brought into the same form and solved using the iterative method introduced in this section.

To account for anisotropy, the scalar permittivity $\epsilon(\boldsymbol{x})$ is replaced by the 3×3 matrix $\boldsymbol{\epsilon}(\boldsymbol{x})$ function in the spatial coordinates, $\boldsymbol{x}$. In this more general case the constitutive relation for the electric displacement field (1) becomes

$$\boldsymbol{D}(\boldsymbol{x}) \quad = \quad \epsilon_0\boldsymbol{\epsilon}(\boldsymbol{x})\boldsymbol{E}(\boldsymbol{x}), \tag{9}$$

and similarly the susceptibility becomes $\boldsymbol{\chi}(\boldsymbol{x}) = \boldsymbol{\epsilon}(\boldsymbol{x}) - \alpha\mathbb{1}_3$. Applying this constitutive relation to the Maxwell equations as we did to obtain (2), the vector Helmholtz equation becomes

$$\boldsymbol{\nabla}\times\boldsymbol{\nabla}\times\boldsymbol{E} - k_0^2\alpha\boldsymbol{E}(\boldsymbol{x}) - k_0^2\boldsymbol{\chi}(\boldsymbol{x})\boldsymbol{E}(\boldsymbol{x}) \quad = \quad \boldsymbol{S}(\boldsymbol{x}). \tag{10}$$

An anisotropic analogue of the modified Born series is given by replacing the scalar preconditioner with a matrix function, $\boldsymbol{\Gamma}(\boldsymbol{x}) \equiv \frac{\mathrm{i}}{\alpha_\mathrm{i}}\boldsymbol{\chi}(\boldsymbol{x})$, in equation (8). The iteration matrix can now be written as a function of the matrix function $\boldsymbol{\chi}(\boldsymbol{x})$:

$$\boldsymbol{M} \quad \equiv \quad \frac{\mathrm{i}k_0^2}{\alpha_\mathrm{i}}\boldsymbol{\chi}\overleftrightarrow{\boldsymbol{G}}\boldsymbol{\chi} - \frac{\mathrm{i}}{\alpha_\mathrm{i}}\boldsymbol{\chi} + \mathbb{1}_3. \tag{11}$$

The question is whether there still exists a choice of α_i such that all the eigenvalues of $\boldsymbol{M}$ have a magnitude less than unity. To show that there is such a choice of α_i we consider the numerical radius, $\max_{\boldsymbol{n}}|\langle\boldsymbol{n}|\boldsymbol{M}|\boldsymbol{n}\rangle|$, where $\langle\boldsymbol{n}|\boldsymbol{n}\rangle = 1$. Evidently the numerical radius will always be at least as large as the largest magnitude eigenvalue of $\boldsymbol{M}$ and thus we can enforce the convergence of the series (8) through the requirement

$$\max_{\boldsymbol{n}}\left|\left\langle\boldsymbol{n}\left|\frac{\mathrm{i}k_0^2}{\alpha_\mathrm{i}}\boldsymbol{\chi}\overleftrightarrow{\boldsymbol{G}}\boldsymbol{\chi} - \frac{\mathrm{i}}{\alpha_\mathrm{i}}\boldsymbol{\chi} + \mathbb{1}_3\right|\boldsymbol{n}\right\rangle\right| \quad < \quad 1 \tag{12}$$

In Supplementary Section S2 it is shown that this requirement can be satisfied by choosing α_i to be larger than $\|\boldsymbol{\Delta}\|$, the largest singular value of $\boldsymbol{\Delta} \equiv \boldsymbol{\epsilon} - \alpha_\mathrm{r}\mathbb{1}_3$, under the conditions that the eigenvectors of $\boldsymbol{\Delta}$ are orthogonal and that the material is free of gain with non-zero losses. Orthogonal eigenvectors are common in practice, e.g. when there is no point with both anisotropic absorption and birefringence in the material. In what follows we further show that a sufficiently large value of α_i can ensure convergence for any material with non-zero losses.

The dyadic Green function, $\overleftrightarrow{\boldsymbol{G}}$, in convergence condition (12) can be written in terms of the identity and an unitary operator $\boldsymbol{U}$ as $\overleftrightarrow{\boldsymbol{G}} \equiv -(2\mathrm{i}k_0^2\alpha_i)^{-1}(\mathbb{1}_3 - \boldsymbol{U})$ (see Supplementary Section S1B). Condition (12) can then be rewritten using the triangle inequality $|a+b| \le |a| + |b|$ as

$$\max_{\boldsymbol{n}}\left[\left|\langle\boldsymbol{n}|\boldsymbol{\Delta}^2 - \alpha_\mathrm{i}^2\mathbb{1}_3|\boldsymbol{n}\rangle\right| + |\langle\boldsymbol{n}\,|\boldsymbol{\chi}\boldsymbol{U}\boldsymbol{\chi}|\,\boldsymbol{n}\rangle|\right] \quad < \quad 2\alpha_\mathrm{i}^2 \tag{13}$$

where $\boldsymbol{\Delta} = \boldsymbol{\chi} + \mathrm{i}\alpha_\mathrm{i}\mathbb{1}_3$. To simplify this condition further we apply the Cauchy-Schwarz inequality $|\langle\boldsymbol{a}|\boldsymbol{b}\rangle|^2 \le \langle\boldsymbol{a}|\boldsymbol{a}\rangle\langle\boldsymbol{b}|\boldsymbol{b}\rangle$ to remove the unitary operator from the second term, followed by the inequality $a^2 + b^2 \ge 2ab$

$$\left|\langle\boldsymbol{n}|\boldsymbol{\Delta}^2 - \alpha_\mathrm{i}^2|\boldsymbol{n}\rangle\right| + \left\langle\boldsymbol{n}\left|\frac{1}{2}(\boldsymbol{\Delta}\boldsymbol{\Delta}^\dagger + \boldsymbol{\Delta}^\dagger\boldsymbol{\Delta})\right|\boldsymbol{n}\right\rangle - 2\alpha_\mathrm{i}\langle\boldsymbol{n}|\boldsymbol{\Delta}_\mathrm{i}|\boldsymbol{n}\rangle < \alpha_\mathrm{i}^2 \quad \forall\boldsymbol{n} : \|\boldsymbol{n}\| = 1, \tag{14}$$

where $\boldsymbol{\Delta} = \boldsymbol{\Delta}_\mathrm{r} + \mathrm{i}\boldsymbol{\Delta}_\mathrm{i}$, with $\boldsymbol{\Delta}_\mathrm{r}$ and $\boldsymbol{\Delta}_\mathrm{i}$ Hermitian matrices which depend respectively on the reactive and dissipative response of the material. As a final step we rewrite $\boldsymbol{\Delta}^2 = (1/2)(\boldsymbol{\Delta}\boldsymbol{\Delta}^\dagger + \boldsymbol{\Delta}^\dagger\boldsymbol{\Delta}) + \mathrm{i}(\boldsymbol{\Delta}\boldsymbol{\Delta}_\mathrm{i} + \boldsymbol{\Delta}_\mathrm{i}\boldsymbol{\Delta})$ and apply the triangle inequality for a second time

$$\left|\left\langle\boldsymbol{n}\left|\frac{1}{2}\left(\boldsymbol{\Delta}\boldsymbol{\Delta}^\dagger + \boldsymbol{\Delta}^\dagger\boldsymbol{\Delta}\right)\right|\boldsymbol{n}\right\rangle - \alpha_\mathrm{i}^2\right| + |\langle\boldsymbol{n}\,|\boldsymbol{\Delta}\boldsymbol{\Delta}_\mathrm{i} + \boldsymbol{\Delta}_\mathrm{i}\boldsymbol{\Delta}|\,\boldsymbol{n}\rangle| + \left\langle\boldsymbol{n}\left|\frac{1}{2}\left(\boldsymbol{\Delta}\boldsymbol{\Delta}^\dagger + \boldsymbol{\Delta}^\dagger\boldsymbol{\Delta}\right)\right|\boldsymbol{n}\right\rangle - 2\alpha_\mathrm{i}\langle\boldsymbol{n}|\boldsymbol{\Delta}_\mathrm{i}|\boldsymbol{n}\rangle < \alpha_\mathrm{i}^2 \quad \forall\boldsymbol{n} : \|\boldsymbol{n}\| = 1. \tag{15}$$

We are now in a position to give a condition for the convergence of our modified Born series. Assuming that $\boldsymbol{\Delta}_\mathrm{i}$ is a positive definite operator everywhere, equivalent to assuming a gain-free material that exhibits a dissipative response throughout space, we can restrict α_i to be greater than the larger of the following two quantities:

$$A_1 = \max_{\boldsymbol{n}}\sqrt{\frac{1}{2}\left\langle\boldsymbol{n}\left|\boldsymbol{\Delta}\boldsymbol{\Delta}^\dagger + \boldsymbol{\Delta}^\dagger\boldsymbol{\Delta}\right|\boldsymbol{n}\right\rangle}, \tag{16}$$

$$A_2 = \max_{\boldsymbol{n}}\frac{|\langle\boldsymbol{n}|\boldsymbol{\Delta}\boldsymbol{\Delta}_\mathrm{i} + \boldsymbol{\Delta}_\mathrm{i}\boldsymbol{\Delta}|\boldsymbol{n}\rangle|}{2\langle\boldsymbol{n}|\boldsymbol{\Delta}_\mathrm{i}|\boldsymbol{n}\rangle}. \tag{17}$$

If the convergence condition

$$\alpha_\mathrm{i} \quad > \quad \max\{A_1, A_2\}, \tag{18}$$

is fulfilled then the inequality (15) will be satisfied, and the modified Born series will converge for the anisotropic medium in question.

For this particular bound it is very important that $\boldsymbol{\Delta}_\mathrm{i}$ is positive definite, rather than non-negative. If $\boldsymbol{\Delta}_\mathrm{i}$ has a kernel containing even one eigenvector $|\boldsymbol{n}_0\rangle$ then for vectors $|\boldsymbol{n}\rangle = |\boldsymbol{n}_0\rangle + \eta|\boldsymbol{n}_\perp\rangle$, A_2 diverges as $\sim 1/\eta$, as $\eta \to 0$. This is because A_2 is the analogue of a weak value[26] of $\boldsymbol{\Delta}_{\mathrm{r,i}}$ with respect to the two vectors $|\boldsymbol{n}\rangle$ and $\boldsymbol{\Delta}_\mathrm{i}|\boldsymbol{n}\rangle$, a value which is well known to potentially lie outside the spectrum of the operator when the two vectors are close to orthogonal (so-called superweak values[27]). In Supplementary Section S2B it is shown that when $\boldsymbol{\Delta}_\mathrm{i}$ has an empty kernel, no such divergence arises, and the Born series can be made to converge for any anisotropic medium. As it stands, it is difficult to make use of condition (18), because the value of A_2 cannot be readily estimated. Extensive numerical tests (Supp. Sec. S2B) have also indicate the larges singular value $\|\boldsymbol{\Delta}\|$ is a tighter bound for α_i that is valid for general gain-free materials, although we have not been able to show this analytically. As a diverging series is straightforward to detect, we have set $\alpha_\mathrm{i} = \|\boldsymbol{\Delta}\|$ in all our simulations. Note that in the limiting case where the medium is isotropic and non-magnetic, $\boldsymbol{\Delta} = \Delta\mathbb{1}_3$, and the basis $|\boldsymbol{n}\rangle$ is the position basis, convergence

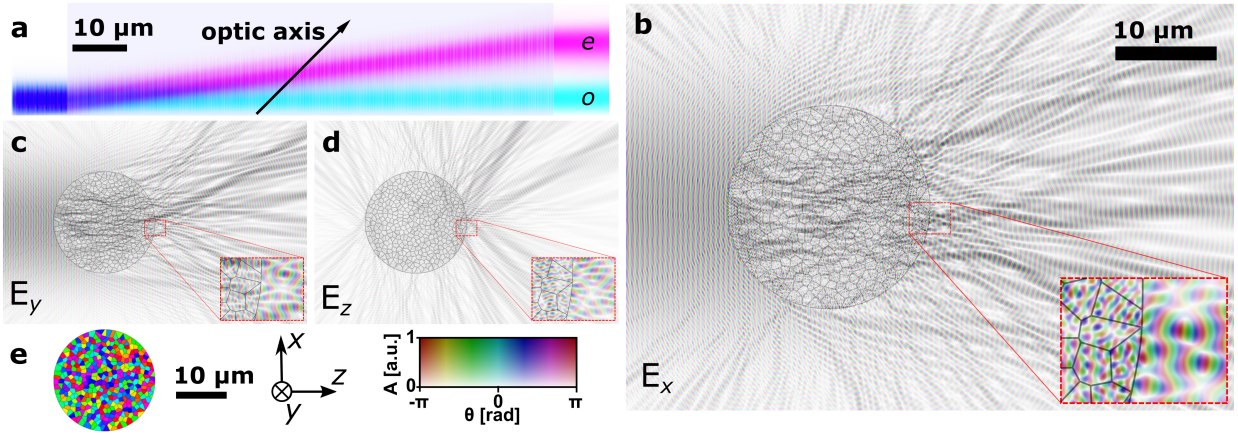

Figure 1: (a) Demonstration of anisotropic permittivity. Diagonally polarized light propagates from left to right through a calcite crystal (light gray box) cut at 45° with respect to its optic axis (indicated by the arrow). It can be seen that, as expected, the in-plane polarized extraordinary ray (e, magenta) is displaced from the ray that is polarized perpendicular to the plane (o, cyan). Some interference can be noticed between the incoming wave and its back-reflection at both the entrance and exit surface of the crystal. (b-e) A circularly polarized Gaussian beam incident from the left on a birefringent vaterite ($CaCO_3$) microrod with a diameter of 20 µm forms a complex scattering pattern instead of a single focus. Although the volume is homogeneous $CaCO_3$, complex, seemingly random, scattering occurs due to its subdivision in crystals of approximately 1 µm in cross-section for which the fast axis is oriented randomly with angles θ shown as hue in panel (e). Panels (b-d) show the field components E_x, E_y, and E_z, respectively. The darkness and hue indicate the field amplitude and phase, respectively, as indicated by the legend in panel (e). An overlaid gray grid outlines the crystal areas for reference, and the inset shows a 4× magnified detail of the field at the exit surface.

condition (18) reduces to $\alpha_\mathrm{i} > \max_{\boldsymbol{x}} |\Delta(\boldsymbol{x})|$. This reproduces the convergence condition found by both Osnabrugge et al.[17] and Krüger et al.[18].

To demonstrate that this series can be used to calculate the propagation of light through an anisotropic material, we consider a birefringent crystal. Perhaps the most common example is calcite ($CaCO_3$, $n_o = 2.776$, $n_e = 2.219$ at $\lambda = 500$ nm), which splits an incident beam into two orthogonally-polarized beams that travel along different paths[28]. Fig. 1a shows that the modified Born series reproduces this effect for a circularly-polarized Gaussian beam (wavelength of 500 nm) traversing the crystal from left to right. The crystal displaces the extra-ordinary polarization laterally (e, magenta), while the ordinary-polarized component (o, cyan) travels along the original optical axis unaffected by the anisotropy.

This method of computing the solution to Maxwell's equation is quite unusual compared to, for instance, a finite difference calculation. Firstly, each iteration can be performed in constant space and involves only a product of diagonal matrices and Fourier transforms, both of which have relatively low computational demands. Secondly, and perhaps even more interesting, the iterative corrections to the field predominantly depend on the local field. This is because the modified Green function, $\overleftrightarrow{\boldsymbol{G}}(\boldsymbol{x}, \boldsymbol{x}')$, is the solution to equation (4), which is that for a dissipative medium characterized by a positive α_i. Both of these features can be advantageous for the efficient simulation of wave propagation in large heterogeneous materials. To illustrate this we simulated ($\lambda = 500$ nm) the electric field within and behind a heterogeneous calcite rod of diameter 20 µm made up of crystal grains with a variable diameter of approximately 1 µm and with a random orientation (Fig. 1b-e). An isotropic material would have focused both polarizations. Instead, a highly irregular speckle pattern can be seen to emerge, elongated along its propagation direction. The insets show close-ups of the high irregularity in the three field-components as the beam exits at the $CaCO_3$-air interface.

In the case of ordinary birefringence, the 3×3-matrix that represents the permittivity at a specific point in space is Hermitian. More general materials may have a permittivity that is non-Hermitian. To illustrate the application of the modified Born series to lossy anisotropic media we calculated the transmission through the simplest such system: an absorbing polarizer. Fig. 2 shows a series of polarizers with varying alignment and an anisotropic extinction coefficient of $\kappa = 0.1$. For clarity, back reflections are avoided by setting the refractive index of the polarizers to the same value as that of the embedding medium ($n = 1$). The two geometries shown are two cross-polarizers (a,c,e), the second of which blocks all transmission. When a third polarizer is inserted in-between the same two polarizers, and with its transmission axis at 45° to the first two polarization axes (b,d,f), this results in a non-zero transmission in agreement with Malus' law.

Algorithm 1 shows the pseudo code to calculate the electric field in anisotropic non-magnetic materials. Before starting the iteration loop, the algorithm must choose a background permittivity, α that will lead to a fast convergence. This is done by numerically finding the value α_r that minimizes the largest singular value of $\boldsymbol{\Delta} = \boldsymbol{\epsilon}(\boldsymbol{x}) - \alpha_\mathrm{r}\mathbb{1}_3$, using the Nelder-Mead algorithm[29]. The imaginary part of the background permittivity, α_i, is then set to equal the minimized upper bound for the singular value $\|\boldsymbol{\Delta}\|$. Although the strict equality does not necessarily guarantee convergence, in practice we have never encountered divergence. It should be noted that calculating the largest singular value is impractical for large problems, so an upper bound must be calculated instead and α_i will likely be larger than the largest singular value, $\|\boldsymbol{\Delta}\|$. For good measure, the complete algorithm (S1) presented in Supplementary Section S1, does include a check for divergence, though in practice it has never been encountered. The algorithms' description is followed by a discussion of its efficiency, robustness to errors, and the issues of sampling and aliasing.

The iteration loop starts on line 6 and repeats until the con-

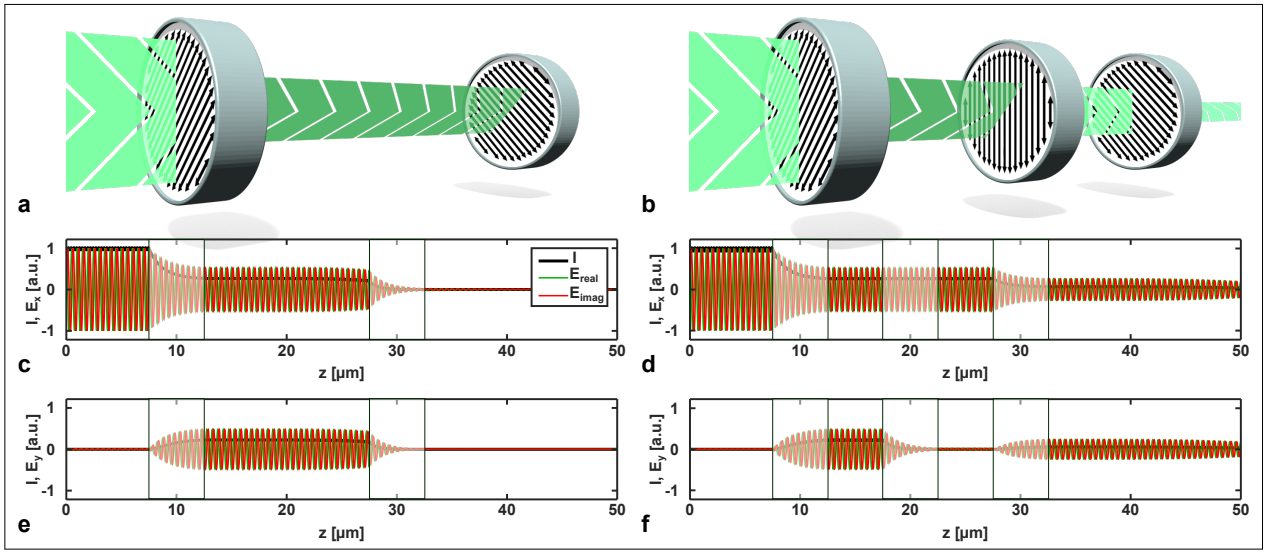

Figure 2: Demonstration of non-hermitian anisotropy. A x-polarized wave traverses two (a,c,e), or three (b,d,f), polarizers from the left to the right. The first and the last polarizer are in cross-diagonal-orientation, preventing transmission through the first system shown in (a). The electric field and intensity are shown for the vertical and horizontal components in (c,e) and (d,f), respectively. Arbitrary units are used so that the maximum value of the intensity and field match for clarity of display.

dition on line 9 is met. In each loop a correction term, $\Delta \boldsymbol{E}$, is calculated for the electric field estimate, $\boldsymbol{E}$. Unless divergence is detected, the correction term is added to the current electric field estimate. In the unlikely event that divergence is detected, the current iteration is repeated for a more conservative, larger, imaginary value for the permittivity bias, α_i. This procedure is repeated until the correction term becomes smaller than a preset tolerance, r_max. Alternative stopping criteria may be considered to improve the algorithm's performance[18].

Algorithm 1 Calculation of the electric field in materials with anisotropic permittivity, $\boldsymbol{\epsilon}$.

1: **function** SOLVEANISOTROPIC($\Delta \boldsymbol{r}, \boldsymbol{j}, \boldsymbol{\epsilon}, \boldsymbol{E_0} = \boldsymbol{0}, r_\mathrm{max}$)
2: $\quad \alpha_\mathrm{r} \leftarrow \underset{\alpha_\mathrm{r}}{\mathrm{argmin}} \left[\| \boldsymbol{\epsilon} - \mathbb{1}_3 \alpha_\mathrm{r} \| \right] \quad \forall \alpha_\mathrm{r} \in \mathbb{R}$
3: $\quad \alpha_\mathrm{i} \leftarrow \| \boldsymbol{\epsilon} - \mathbb{1}_3 \alpha_\mathrm{r} \|$
4: $\quad \boldsymbol{\chi} \leftarrow \boldsymbol{\epsilon} - (\alpha_\mathrm{r} + \mathrm{i}\alpha_\mathrm{i})\, \mathbb{1}_3$
5: $\quad \boldsymbol{E} \leftarrow \boldsymbol{E_0}$
6: $\quad$ **repeat**
7: $\quad\quad \Delta \boldsymbol{E} \leftarrow \frac{\mathrm{i}}{\alpha_\mathrm{i}} \boldsymbol{\chi} \left[\overleftrightarrow{\boldsymbol{G}} * \left(k_0^2 \boldsymbol{\chi} \boldsymbol{E} + \boldsymbol{S} \right) - \boldsymbol{E} \right]$
8: $\quad\quad \boldsymbol{E} \leftarrow \boldsymbol{E} + \Delta \boldsymbol{E}$
9: $\quad$ **until** $\| \Delta \boldsymbol{E} \| < r_\mathrm{max} \| \boldsymbol{E} \|$
10: $\quad$ **return** $\boldsymbol{E}$
11: **end function**

3 Magnetic and chiral materials

As it stands, Algorithm 1 can only be used on non-magnetic materials. However, the preconditioning step can be generalized to account for general magnetic media, without requiring changes to the iteration loop. A general algorithm is given in Supplementary Section S1A. The susceptibility, $\boldsymbol{\chi}$, must be redefined for magnetic materials to account for permeability and any electro-magnetic coupling constants. A permeability scale, β, is introduced, and the background permittivity, α, is determined from the modified susceptibility.

The constitutive relations for general linear electromagnetic materials are given by (see e.g.[30]):

$$\begin{aligned} \boldsymbol{D}(\boldsymbol{x}) &= \epsilon_0 \boldsymbol{\epsilon}(\boldsymbol{x}) \boldsymbol{E}(\boldsymbol{x}) + \frac{1}{c} \boldsymbol{\xi}(\boldsymbol{x}) \boldsymbol{H}(\boldsymbol{x}) \\ \boldsymbol{B}(\boldsymbol{x}) &= \mu_0 \boldsymbol{\mu}(\boldsymbol{x}) \boldsymbol{H}(\boldsymbol{x}) + \frac{1}{c} \boldsymbol{\zeta}(\boldsymbol{x}) \boldsymbol{E}(\boldsymbol{x}). \end{aligned} \tag{19}$$

Here both the permittivity $\boldsymbol{\epsilon}(\boldsymbol{x})$ and permeability, $\boldsymbol{\mu}(\boldsymbol{x})$, are tensors that depend on the position, $\boldsymbol{x}$. The additional coupling terms $\boldsymbol{\xi}(\boldsymbol{x})$ and $\boldsymbol{\zeta}(\boldsymbol{x})$ enable arbitrary linear interactions between the electric and magnetic components, which are commonplace in—for instance—metamaterials[31]. The coupling tensors $\boldsymbol{\xi}(\boldsymbol{x})$ and $\boldsymbol{\zeta}(\boldsymbol{x})$ are also essential to model chiral materials[32], and non-reciprocal materials such as moving media[33], and Tellegen media[25].

Substituting (19) into the Maxwell equations and eliminating the magnetic field, $\boldsymbol{H} = (\mathrm{i}\omega\mu_0)^{-1} \boldsymbol{\mu}^{-1} [\boldsymbol{\nabla} \times \boldsymbol{E} - \mathrm{i} k_0 \boldsymbol{\zeta} \boldsymbol{E}]$, leads to the same form of vector Helmholtz equation (10). It is sufficient to generalize the source and susceptibility as follows:

$$\boldsymbol{S} \equiv \mathrm{i}\omega\mu_0 \beta^{-1} \boldsymbol{j} \tag{20}$$

$$\begin{aligned} \boldsymbol{\chi} = \; & \frac{\boldsymbol{\epsilon}}{\beta} - \frac{1}{\beta} \boldsymbol{\xi} \boldsymbol{\mu}^{-1} \boldsymbol{\zeta} - \mathbb{1}_3 \alpha + - \frac{\mathrm{i}}{\beta} \boldsymbol{\xi} \boldsymbol{\mu}^{-1} \boldsymbol{\mathcal{D}} \\ & + \frac{\mathrm{i}}{\beta} \boldsymbol{\mathcal{D}} \boldsymbol{\mu}^{-1} \boldsymbol{\zeta} + \boldsymbol{\mathcal{D}} \left(\mathbb{1}_3 - \frac{\boldsymbol{\mu}^{-1}}{\beta} \right) \boldsymbol{\mathcal{D}}, \end{aligned} \tag{21}$$

where $\boldsymbol{\mathcal{D}} \equiv k_0^{-1} \boldsymbol{\nabla} \times$ is the differential operator, which calculation is discussed in Supplementary Section S1A. For conciseness, the spatial dependency of the variables is omitted in equations (20) and (21). Note that all these operations can be implemented efficiently using fast-Fourier transforms and diagonal-matrix multiplications. Since vector Helmholtz equation (10) still applies, the same iterative procedure can be used to solve for non-magnetic anisotropic materials.

Although the modified Born series is now formally identical for these general constitutive relations (19), to that for the anisotropic dielectric discussed in section 2, it is important to note that the susceptibility (21) is no longer a diagonal matrix in the position variable. This is because it now contains the differential operators $\boldsymbol{\mathcal{D}}$. Although this affects neither the form of the modified Born series (i.e. as a product of diagonal matrices and

Fourier transforms), or the proof of its convergence (which did not make any assumptions about the form of $\boldsymbol{\chi}$), it can make the convergence of the series less rapid. This is because the largest singular value of $\boldsymbol{\mathcal{D}}$ has a magnitude of $\pi/(k_0\Delta r)$, where Δr is the grid spacing. Therefore, the required number of iterations increases with the density of the sampling grid when simulating general magnetic materials.

The optimal value of α is determined using a similar method in both the all-dielectric and general cases. Albeit not critical for convergence, the additional parameter β we have introduced in (21) can be used to enable further optimization of the convergence speed.

Algorithm S1 in Supplementary Section S1A generalizes Algorithm 1 to magnetic materials by redefining $\hat{\sigma}$ as a function of two variables, α_r and β, and minimizing it in both variables. Using the usual inequalities for matrix norms[34], the function $\hat{\sigma}(\beta, \alpha_\mathrm{r})$ determines an upper bound for the singular values of the separate terms in the sum (21), and using the largest singular value, $\sigma_{\mathcal{D}} = \pi/(k_0\Delta r)$, of the discretized differential operator, $\boldsymbol{\mathcal{D}}$. The Nelder-Mead method is used to numerically minimize the product $\hat{\sigma}(\beta, \alpha)\beta$, which we found to correlate well with the inverse of the convergence rate. The minimization provides values for the permeability scale, β, and the real part of the background permittivity α_r. The imaginary part, α_i, is set to $\hat{\sigma}(\beta, \alpha)$, an upper bound estimate for the largest singular value of $\boldsymbol{\Delta} = \boldsymbol{\chi} + \mathrm{i}\alpha_\mathrm{i}\mathbb{1}_3$.

At optical frequencies, it is a common approximation to assume that the permeability, $\boldsymbol{\mu}$, equals $\mu_0\mathbb{1}_3$ in the entire simulation volume. While this is appropriate at such frequencies (where magnetism can in many circumstances be ignored[33,35]), the modified Born series can—in principle—be applied to any electromagnetic problem, where magnetic effects are important. Indeed, while all our simulations are apparently for optical length scales we must remember that Maxwell's equations are scale invariant so that the same results would be obtained with the spatial frequencies, $\boldsymbol{k}$, scaled by the same factor as the temporal frequency, ω. In Supplementary Figure S4(a-f), we show that our generalized Born series can simulate impedance matched media (where $\boldsymbol{\epsilon} = \boldsymbol{\mu}$, which includes transformation media[36], as well as other interesting wave effects (e.g.[37]).

Since the constitutive relations (19) include the electric-magnetic coupling tensors, $\boldsymbol{\zeta}$ and $\boldsymbol{\xi}$, one can also use the modified Born series to treat bi-isotropic and more generally bi-anisotropic media. Many organic molecules such as glucose have a chiral asymmetry that leads to bi-isotropy. This chirality causes a rotation of linear polarization around the axis of propagation. As a second example of our generalized approach, Supplementary Figures S4(g and h) show simulations where linearly polarized light is slowly rotated in a chiral solution. As the rotation takes several millimeters to complete, the calculation was performed for a wavelength of 500 nm and over a propagation length of 10 mm.

As a final example we consider negative index materials. Simultaneously negative values for $\boldsymbol{\epsilon}$ and $\boldsymbol{\mu}$ can be obtained by engineering a material at the sub-wavelength scale. Such materials are characterized by a negative refractive index[38,39]. The general algorithm presented here naturally handles negative values for both the permittivity, $\boldsymbol{\epsilon}$, and the permeability, $\boldsymbol{\mu}$. Fig. 3 demonstrates this for a block of transparent material, yet with a negative refractive index of -1.5 ($\epsilon = -1.5$, $\mu = -1$). A Gaussian beam is incident at 30° to the surface normal and refracts backwards into the material with an angle of $-19.5°$, opposite to what is expected from Snell's law for regular glass. The wave appears to travel backwards from left to right inside the metamaterial. At the exit surface the beam couples out at 30°, again opposing Snell's law. Because there is no impedance matching in this case, some interference can be seen between the incident and reflected beam in Fig. 3.

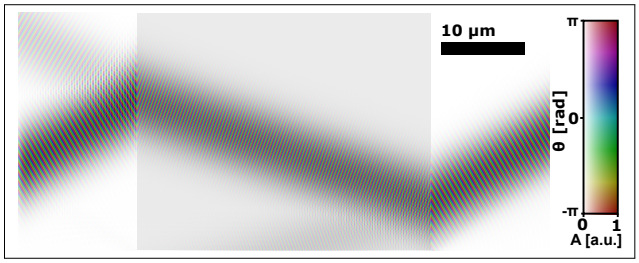

Figure 3: Light-wave propagation through a material with negative refractive index ($\epsilon = -1.5$, $\mu = -1$, gray). The field amplitude (brightness) and phase (hue) are shown for a Gaussian beam that enters the surface at 30° to the normal. The beam can be seen to refract at $-19.5°$, backwards, into the metamaterial at an angle opposite to that for normal glass.

4 Conclusion

We have extended the computationally efficient modified Born series[17,18], to arbitrary linear electromagnetic media. This enables the accurate calculation of light propagation through birefringent, chiral, and even magnetic materials. We have implemented this numerically, and shown that a large variety of electromagnetic effects can be simulated using this method. The complete is given in Algorithm S1, and implemented as a numerical library that can be found with working examples at `https://uod.box.com/v/macromax`. While there is no doubt that finite element simulations continue to be a very versatile and effective to invert the Helmholtz operator, the modified Born series has features that may make it more suitable for simulating complex samples as those found in microscopy. As the finite element method, it calculates the field due to a fixed frequency source; however, it does this through computing a series of terms that iteratively propagate the field out from the source, not entirely unlike a time-dependent calculation such as FDTD (see for example[17]). Yet, unlike an FDTD calculation, the solution at a given iteration is not considered a boundary condition for the next iteration, rather an estimate of the solution to Maxwell's equations that is improved with each iteration. As such, numerical errors do not accumulate over time. Moreover, this opens the door to precalculate approximate solutions using, for instance, the beam propagation method (Supp. Section S1E). The modified Born series—as it was originally intended[17]—is best suited for large scale problems with a limited range of susceptibilities. The modest memory requirements of the anisotropic modified Born iteration make it an ideal candidate to study light propagation in large heterogeneous tissues or materials (Supp. Section S1C). Our extended method allows for a wide range of material properties, including birefringence, polarization, variable permeability, chirality, and negative refractive index. This enables it to account for the birefringence and chiral effects of light traversing biological tissue. The algorithm is also expected to find use in scattering studies, where laboratory experiments are often performed on highly scattering powders that typically consist of birefringent particles such as rutile powder (TiO_2). The presented numerical method offers a bridge between approximate analytic predictions and experiments. In addition, such large-scale problems are not exclusive to optics and occur across wave physics, including problems involving metamaterials, to which the modified Born series can now be applied.

Funding Information

This work is supported by the Philip Leverhulme Prize from the Leverhulme Trust. SARH acknowledges financial support from the Royal Society and TATA (RPG-2016-186).

References

[1] O. G. Ernst and M. J. Gander. *Numerical Analysis of Multiscale Problems.*, chapter Why it is Difficult to Solve Helmholtz Problems with Classical Iterative Methods, pages 325–363. Springer, Berlin, Heidelberg, 2012. ISBN 978-3-642-22060-9, Online ISBN 978-3-642-22061-6.

[2] J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade. *Photonic Crystals: Molding the Flow of Light - Second Edition*. Princeton University Press, March 2008. ISBN-10: 0691124566, ISBN-13: 9780691124568.

[3] H. C. van de Hulst. *Light Scattering by Small Particles*. Dover Publications Inc., March 2003. ISBN-10: 0486642283, ISBN-13: 978-0486642284.

[4] Stefan Rotter and Sylvain Gigan. Light fields in complex media: Mesoscopic scattering meets wave control. *Rev. Mod. Phys.*, 89:015005, March 2017.

[5] I. M. Vellekoop and A. P. Mosk. Focusing coherent light through opaque strongly scattering media. *Optics Letters*, 32(16):2309, 8 2007.

[6] I. M. Vellekoop, A. Lagendijk, and A. P. Mosk. Exploiting disorder for perfect focusing. *Nature Photon.*, 4:320–322, 2010.

[7] Adam K Glaser, Ye Chen, and Jonathan TC Liu. Fractal propagation method enables realistic optical microscopy simulations in biological tissues. *Optica*, 3(8):861–869, 2016.

[8] Zdzislaw Salamon and Gordon Tollin. Optical anisotropy in lipid bilayer membranes: coupled plasmon-waveguide resonance measurements of molecular orientation, polarizability, and shape. *Biophysical journal*, 80(3):1557–1567, 2001.

[9] MJ Everett, K Schoenenberger, BW Colston, and LB Da Silva. Birefringence characterization of biological tissue by use of optical coherence tomography. *Optics letters*, 23(3):228–230, 1998.

[10] Wen-Chuan Kuo, Shey-Chien Lin, and Chung-Yu Chuang. Birefringence measurement in polarization-sensitive optical coherence tomography using differential-envelope detection method. *Review of Scientific Instruments*, 81(5):053705, 2010.

[11] Gilbert Jarry, Florence Henry, and Robin Kaiser. Anisotropy and multiple scattering in thick mammalian tissues. *JOSA A*, 17(1):149–153, 2000.

[12] Johannes F De Boer, Thomas E Milner, Martin JC van Gemert, and J Stuart Nelson. Two-dimensional birefringence imaging in biological tissue by polarization-sensitive optical coherence tomography. *Optics letters*, 22(12):934–936, 1997.

[13] Sanaz Alali, Karen J Aitken, Annette Schröder, Darius J Bagli, and I Alex Vitkin. Optical assessment of tissue anisotropy in ex vivo distended rat bladders. *Journal of biomedical optics*, 17(8):0860101–0860108, 2012.

[14] Angelo Pierangelo, André Nazac, Abdelali Benali, Pierre Validire, Henri Cohen, Tatiana Novikova, Bicher Haj Ibrahim, Sandeep Manhas, Clément Fallet, Maria-Rosaria Antonelli, and Antonello-De Martino. Polarimetric imaging of uterine cervix: a case study. *Opt. Express*, 21(12):14120–14130, June 2013.

[15] M. Koike-Tani, T. Tani, S. B. Mehta, A. Verma, and R. Oldenbourgh. Polarized light microscopy in reproductive and developmental biology. *Mol. Reprod. Dev.*, 82(7-8):548–562, August 2013.

[16] David C. Adams, Lida P. Hariri, Alyssa J. Miller, Yan Wang, Josalyn L. Cho, Martin Villiger, Jasmin A. Holz, Margit V. Szabari, Daniel L. Hamilos, R. Scott Harris, Jason W. Griffith, Brett E. Bouma, Andrew D. Luster, Benjamin D. Medoff, and Melissa J. Suter. Birefringence microscopy platform for assessing airway smooth muscle structure and function in vivo. *Science Translational Medicine*, 8(359):359ra131–359ra131, 2016.

[17] Gerwin Osnabrugge, Saroch Leedumrongwatthanakun, and Ivo M Vellekoop. A convergent born series for solving the inhomogeneous helmholtz equation in arbitrarily large media. *Journal of computational physics*, 322:113–124, 2016.

[18] Benjamin Krüger, Thomas Brenner, and Alwin Kienle. Solution of the inhomogeneous maxwell's equations using a born series. *Opt. Express*, 25(21):25165–25182, October 2017.

[19] I. Koutromanos. *Fundamentals of Finite Element Analysis: Linear Finite Element Analysis*. John Wiley & Sons, 2018.

[20] M. Born. Quantenmechanik der stoßvorgänge. *Z. Phys.*, 38(38):803, 1926.

[21] E. Akkermans and G. Montambaux. *Mesoscopic Physics of Electrons and Photons*, chapter Coherent backscattering of light. Cambridge University Press, 2007.

[22] R. G. Newton. *Scattering Theory of Waves and Particles*. Springer, 2013.

[23] S. A. R. Horsley, M. Artoni, and G. C. La Rocca. Spatial kramers-kronig relations and the reflection of waves. *Nature Photon.*, 9:436, June 2015.

[24] J. S. Walker. *Fast Fourier Transforms*. CRC Press, 2017.

[25] B. D. H. Tellegen. The gyrator, a new electric network element. *Philips. Res. Rep.*, 3(2):81–101, April 1948.

[26] Y. Aharonov, D. Z. Albert, and L. Vaidman. How the result of a measurement of a component of the spin of a spin-1/2 particle can turn out to be 100. *Phys. Rev. Lett.*, 60:1351, 1988.

[27] M. V. Berry and P. Shukla. Typical weak and superweak values. *J. Phys. A*, 43(43):354024, August 2010.

[28] G. Waldman. *Introduction to Light: The Physics of Light, Vision, and Color*. Dover, New York, 2002.

[29] J. A. Nelder and R. Mead. A simplex method for function minimization. *The Computer Journal*, 7:308–313, 1965.

[30] A. Priou, A. Sihvola, S. Tretyakov, and A. Vinogradov. *Advances in Complex Electromagnetic Materials*. Springer, 2012.

[31] D. R. Smith. Calculation and measurement of bianisotropy in a split ring resonator metamaterial. *J. Appl. Phys.*, 100:024507, 2006.

[32] A. Lakhtakia. *Beltrami Fields in Chiral Media*. World Scientific, 1994.

[33] L. D. Landau, L. P. Pitaevskii, and E. M. Lifshitz. *Electrodynamics of Continuous Media*. Butterworth-Heinemann, 2004.

[34] I. S. Gradsheyn and I. M. Ryzhik. *Table of Integrals Series and Products*. Academic Press, 2000.

[35] R. Merlin. Metamaterials and the landau–lifshitz permeability argument: Large permittivity begets high-frequency magnetism. *Proc. Natl. Acad. Sci. U.S.A.*, 106:1693, 2009.

[36] J. B. Pendry, D. Schurig, and D. R. Smith. Controlling electromagnetic fields. *Science*, 312:1780, 2006.

[37] B. Vial, Y. Liu, S. A. R. Horsley, T. G. Philbin, and Y. Hao. A class of invisible inhomogeneous media and the control of electromagnetic waves. *Phys. Rev. B*, 94:245119, 2016.

[38] V. G. Veselago. The electrodynamics of substances with simultaneously negative values of ϵ and μ. *Sov. Phys. Usp.*, 10:509, 1967.

[39] J. B. Pendry. Negative refraction makes a perfect lens.

Phys. Rev. Lett., 85:3966, 2000.

Supplementary Information: Calculating coherent light-wave propagation in large heterogeneous media.

T. Vettenburg[1,2,*], S.A.R. Horsley[1], and J. Bertolotti[1]

[1] *Department of Physics and Astronomy, University of Exeter, Exeter, EX4 4QL, United Kingdom*
[2] *School of Science and Engineering, University of Dundee, Nethergate, Dundee, DD1 4HN, United Kingdom*
[*] *Corresponding author: t.vettenburg@dundee.ac.uk*

Abstract

Understanding the interaction of light with a highly scattering material is essential for optical microscopy of optically thick and heterogeneous biological tissues. Ensemble-averaged analytic solutions cannot provide more than general predictions for relatively simple cases. Yet, biological tissues contain chiral organic molecules and many of the cells' structures are birefringent, a property exploited by polarization microscopy for label-free imaging. Solving Maxwell's equations in such materials is a notoriously hard problem. Here we present an efficient method to determine the propagation of electro-magnetic waves in arbitrary anisotropic materials. We demonstrate how the algorithm enables large scale calculations of the scattered light field in complex birefringent materials, chiral media, and even materials with a negative refractive index.

S1 The algorithm

S1A Detailed description of the general algorithm

The iteration loop of the algorithm is relatively short and identical for magnetic and non-magnetic materials (Algorithm S1, lines 15-23). Only the preconditioning steps differ. The algorithm starts by checking whether the material has magnetic properties and it determines the background permittivity, $\alpha = \alpha_r + i\alpha_i$, and if required, the permeability scale, β. This enables the definition of the susceptibility, χ, the dyadic Green function, $\overleftrightarrow{G}$, and the source distribution, $\boldsymbol{S}$.

Before starting the iteration loop, the electric field is either initialized to all zero, or to an approximate solution if available. On line 16, the loop starts by calculating the next term in the series, $\Delta\boldsymbol{E}$, using the operators χ and $\overleftrightarrow{G}*$ on the source distribution, $\boldsymbol{S}$, and the current estimate of the field, $\boldsymbol{E}$. The next term is added to the current estimate, $\boldsymbol{E}$, under the condition that the l_2-norm of the new term is less than that of the previous term. Otherwise, the series must be divergent for the current choice of the background permittivity, α, so its imaginary part is increased by 50%. The iteration continues until the updates to the field are deemed sufficiently small.

The susceptibility is defined by $\chi \equiv \left(\epsilon - \xi\mu^{-1}\zeta\right)/\beta - \mathbb{1}_3\alpha - i\xi\mu^{-1}\mathcal{D}/\beta + i\mathcal{D}\mu^{-1}\zeta/\beta + \mathcal{D}\left(\mathbb{1}_3 - \mu^{-1}/\beta\right)\mathcal{D}$, where $\mathcal{D} \equiv k_0^{-1}\nabla\times = k_0^{-1}\mathcal{F}^{-1}\boldsymbol{k}\otimes\mathcal{F}$, with $\otimes$ the outer product, while the forward and inverse Fourier transforms are represented by $\mathcal{F}$ and $\mathcal{F}^{-1}$, respectively. Although all operations can be represented as large matrix operations, it is more space and time efficient to use fast-Fourier transforms and point-wise 3×3-dot products for each application of the operator χ on the electric field, $\boldsymbol{E}$. The source is defined by $\boldsymbol{S} \equiv i\omega\mu_0\ \boldsymbol{j}/\left(\beta k_0^2\right)$. The dyadic Green's function, $\overleftrightarrow{G}*$, is discussed in section S1B in what follows. Note that to prevent issues with numerical precision, in the implementation the factor k_0^{-2} is moved from the definition of $\overleftrightarrow{G}$ to that of $\boldsymbol{S}$.

S1B The dyadic Green's function

The dyadic Green function, $\overleftrightarrow{G}$, is integral to the calculation of the modified Born series:

$$\boldsymbol{E} = \left[\sum_{p=0}^{\infty}\boldsymbol{M}^p\right]\boldsymbol{\Gamma}\overleftrightarrow{G}\boldsymbol{S}, \text{ where} \tag{S1}$$

$$\boldsymbol{M} \equiv k_0^2\boldsymbol{\Gamma}\overleftrightarrow{G}\chi - \boldsymbol{\Gamma} + \mathbb{1}_3, \text{ and} \tag{S2}$$

$$\boldsymbol{\Gamma} \equiv \frac{i}{\alpha_i}\chi. \tag{S3}$$

Although well established in the literature of classical electromagnetism[1], the form of the dyadic Green function that we have presented in equation (6) of the main text may not be familiar to the reader. In this appendix we justify this equation, and show a useful representation of the Green function in terms of a unitary operator, $\boldsymbol{U}$.

The vector Helmholtz equation can be separated into a homogeneous, $\mathcal{O}_h$, and an inhomogeneous part, $\mathcal{O}_i$, as follows

$$\begin{aligned}\nabla\times\nabla\times\boldsymbol{E}(\boldsymbol{x}) - k_0^2\epsilon(\boldsymbol{x})\boldsymbol{E}(\boldsymbol{x}) &= i\omega\mu_0\boldsymbol{j}(\boldsymbol{x})\\ \left(\nabla\times\nabla\times - \alpha k_0^2\right)\boldsymbol{E} - k_0^2(\epsilon-\alpha)\boldsymbol{E} &= \boldsymbol{S}\\ \left(\mathcal{O}_h + \mathcal{O}_i\right)\boldsymbol{E} &= \boldsymbol{S}\end{aligned} \tag{S4}$$

The dyadic Green function, $\overleftrightarrow{G}$, is defined as the impulse response solution to the homogeneous part:

$$\mathcal{O}_h\overleftrightarrow{G} = \nabla\times\nabla\times\overleftrightarrow{G}(\boldsymbol{x},\boldsymbol{x}') - k_0^2\alpha\overleftrightarrow{G}(\boldsymbol{x},\boldsymbol{x}') = \mathbb{1}_3\delta^{(3)}(\boldsymbol{x}-\boldsymbol{x}') \tag{S5}$$

First, we take the Fourier transform of equation (S5), recognizing that—due to the homogeneity of the medium—$\overleftrightarrow{G}$ must be a function of $\boldsymbol{x}-\boldsymbol{x}'$, and thus of a single variable $\boldsymbol{k}$ in Fourier space

$$\overleftrightarrow{G}(\boldsymbol{x},\boldsymbol{x}') = \int\frac{d^3\boldsymbol{k}}{(2\pi)^3}\tilde{\boldsymbol{G}}(\boldsymbol{k})e^{i\boldsymbol{k}(\boldsymbol{x}-\boldsymbol{x}')} \tag{S6}$$

or equivalently in the operator notation used in the main text (where integrals are subsumed into our '' product)

$$\overleftrightarrow{G} = \mathcal{F}^{-1}\tilde{\boldsymbol{G}}\mathcal{F}, \tag{S7}$$

Algorithm S1 A function that implements the general algorithm for both non-magnetic and magnetic materials.

function SOLVEMACROSCOPICMAXWELL($\Delta \boldsymbol{r}, \boldsymbol{j}, \boldsymbol{\epsilon}, \boldsymbol{\xi} = \mathbf{0}, \boldsymbol{\zeta} = \mathbf{0}, \boldsymbol{\mu} = \mathbb{1}_3, \boldsymbol{E_0} = \mathbf{0}, r_{\max}$)
  if $\boldsymbol{\xi} \equiv \boldsymbol{\zeta} \equiv \mathbf{0}$ and $\boldsymbol{\mu}$ is both isotropic and constant **then** ▷ *non-magnetic*
    $\hat{\sigma}(\beta, \alpha_r) \equiv \|\boldsymbol{\epsilon}/\beta - \mathbb{1}_3 \alpha_r\|$ ▷ $\|\|$ *is defined as the largest singular value.*
    $\beta \leftarrow 1/\boldsymbol{\mu}_{11}$
    $\alpha_r \leftarrow \underset{\alpha_r}{\operatorname{argmin}} \left[\hat{\sigma}(\beta, \alpha_r)\right] \quad \forall \alpha_r \in \mathbb{R}$
  else ▷ *magnetic*
    $\hat{\sigma}(\beta, \alpha_r) \equiv \left\|\left(\boldsymbol{\epsilon} - \boldsymbol{\xi}\boldsymbol{\mu}^{-1}\boldsymbol{\zeta}\right)/\beta - \mathbb{1}_3 \alpha_r\right\| + \left[\left\|\boldsymbol{\xi}\boldsymbol{\mu}^{-1}\right\| + \left\|\boldsymbol{\mu}^{-1}\boldsymbol{\zeta}\right\|\right] \sigma_{\mathcal{D}}/\beta + \left\|\mathbb{1}_3 - \boldsymbol{\mu}^{-1}/\beta\right\| \sigma_{\mathcal{D}}^2$
    ▷ *where* $\sigma_{\mathcal{D}} = k_0^{-1}\pi/\|\Delta \boldsymbol{r}\|$ *is proportional to the highest possible spatial frequency*
    $\beta, \alpha_r \leftarrow \underset{\beta, \alpha_r}{\operatorname{argmin}} \left[\hat{\sigma}(\beta, \alpha_r)\,|\beta|\right] \quad \forall \beta \in \mathbb{R}_{>0}, \alpha_r \in \mathbb{R}$
  end if
  $\alpha_i \leftarrow \hat{\sigma}(\beta, \alpha_r)$
  $\alpha \leftarrow \alpha_r + \mathrm{i}\alpha_i$
  $\boldsymbol{\chi} \leftarrow \frac{1}{\beta}\left(\boldsymbol{\epsilon} - \boldsymbol{\xi}\boldsymbol{\mu}^{-1}\boldsymbol{\zeta}\right) - \mathbb{1}_3\alpha - \frac{\mathrm{i}}{\beta}\boldsymbol{\xi}\boldsymbol{\mu}^{-1}\boldsymbol{\mathcal{D}} + \frac{\mathrm{i}}{\beta}\boldsymbol{\mathcal{D}}\boldsymbol{\mu}^{-1}\boldsymbol{\zeta} + \boldsymbol{\mathcal{D}}\left(\mathbb{1}_3 - \frac{1}{\beta}\boldsymbol{\mu}^{-1}\right)\boldsymbol{\mathcal{D}}$
  $\boldsymbol{E} \leftarrow \boldsymbol{E_0}$
  $p \leftarrow \infty$ ▷ *the* l_2*-norm of the previous update*
  repeat
    $\Delta \boldsymbol{E} \leftarrow \frac{\mathrm{i}}{\alpha_i}\boldsymbol{\chi}\left[\overleftrightarrow{\boldsymbol{G}} * \left(k_0^2 \boldsymbol{\chi}\boldsymbol{E} + \boldsymbol{S}\right) - \boldsymbol{E}\right]$ ▷ *calculate the next term in the series*
    if $\|\Delta \boldsymbol{E}\| < p$ **then**
      $\boldsymbol{E} \leftarrow \boldsymbol{E} + \Delta \boldsymbol{E}$ ▷ *update current field estimate*
      $p \leftarrow \|\Delta \boldsymbol{E}\|$
    else
      $\alpha_i \leftarrow 1.5\alpha_i$ ▷ *increase* α_i *if divergence would be detected*
    end if
  until $\|\Delta \boldsymbol{E}\| < r_{\max}\|\boldsymbol{E}\|$
  return $\boldsymbol{E}$
end function

where $\tilde{\boldsymbol{G}}$ is diagonal in Fourier space. Substituting expression (S6) into (S5) we obtain the following equation for the Fourier components of the Green function

$$\boldsymbol{k} \times \boldsymbol{k} \times \tilde{\boldsymbol{G}}(\boldsymbol{k}) + k_0^2 \alpha \tilde{\boldsymbol{G}}(\boldsymbol{k}) = -\mathbb{1}_3. \tag{S8}$$

At this point we decompose the identity matrix on the right into two parts $\mathbb{1}_3 = (\mathbb{1}_3 - \boldsymbol{\Pi}_L) + \boldsymbol{\Pi}_L = \boldsymbol{\Pi}_T + \boldsymbol{\Pi}_L$, where

$$\boldsymbol{\Pi}_L = \frac{\boldsymbol{k} \otimes \boldsymbol{k}}{\boldsymbol{k}^2}. \tag{S9}$$

These two operators can be physically understood as projecting out the longitudinal ($\boldsymbol{\Pi}_L$) and transverse ($\boldsymbol{\Pi}_T$) parts of the electromagnetic field, associated with electrostatic and radiative contributions respectively. We similarly decompose the Green function as $\tilde{\boldsymbol{G}}(\boldsymbol{k}) = g_L(\boldsymbol{k})\boldsymbol{\Pi}_L + g_T(\boldsymbol{k})\boldsymbol{\Pi}_T$, finding that equation (S8) separates into two parts

$$\begin{aligned} k_0^2 \alpha g_L(\boldsymbol{k}) &= -1 \\ \left(\boldsymbol{k}^2 - \alpha k_0^2\right) g_T(\boldsymbol{k}) &= 1 \end{aligned} \tag{S10}$$

from which we can deduce that the Green function is given by

$$\overleftrightarrow{\boldsymbol{G}}(\boldsymbol{x}, \boldsymbol{x}') = \int \frac{d^3\boldsymbol{k}}{(2\pi)^3}\left[\frac{\boldsymbol{\Pi}_T}{\boldsymbol{k}^2 - \alpha k_0^2} - \frac{\boldsymbol{\Pi}_L}{\alpha k_0^2}\right] \mathrm{e}^{\mathrm{i}\boldsymbol{k}(\boldsymbol{x} - \boldsymbol{x}')} \tag{S11}$$

which in our operator notation we write as

$$\overleftrightarrow{\boldsymbol{G}} = \mathcal{F}^{-1}\left(\frac{\boldsymbol{\Pi}_T}{\boldsymbol{k}^2 - \alpha k_0^2} - \frac{\boldsymbol{\Pi}_L}{\alpha k_0^2}\right)\mathcal{F} \tag{S12}$$

where the quantity in the rounded brackets must be understood as a matrix in both vector and Fourier indices. The dependence of the bracketed quantity on a single Fourier space variable $\boldsymbol{k}$ indicates that this matrix is diagonal in the Fourier indices, with each diagonal entry corresponding to a different value of $\boldsymbol{k}$. This completes the demonstration of equation (S12).

The Green function operator (S12) can be written as a linear combination of the identity operator and a unitary operator, an observation that is instrumental in the study of the modified Born series (S1) and its convergence. This representation of the Green function relies on the background permittivity α being chosen as a complex number. To find this representation we first note the following identity for any complex number z

$$\frac{1}{z} = \frac{1}{2\mathrm{i}\,\mathrm{Im}[z]}\left(1 - \frac{z^\star}{z}\right) = \frac{1}{2\mathrm{i}\,\mathrm{Im}[z]}\left(1 - \mathrm{e}^{-2\mathrm{i}\angle z}\right), \tag{S13}$$

where $\angle z$ indicates the complex argument of z. Applying identity (S13) to (S12), the Green function operator (S12) is reduced to the following form

$$\overleftrightarrow{\boldsymbol{G}} = -\frac{1}{2\mathrm{i}\alpha_i k_0^2}\left(\mathbb{1}_3 - \boldsymbol{U}\right) \tag{S14}$$

where α_i is the imaginary part of α and the unitary matrix $\boldsymbol{U}$ is given by

$$\boldsymbol{U} = \mathcal{F}^{-1}\left(\frac{k^2 - \alpha^\star k_0^2}{k^2 - \alpha k_0^2}\boldsymbol{\Pi}_T + \frac{\alpha^\star}{\alpha}\boldsymbol{\Pi}_L\right)\mathcal{F}, \tag{S15}$$

where $k = \|\boldsymbol{k}\|$ is the l_2-norm of $\boldsymbol{k}$, while α^* is the complex conjugate of α. The operator (S15) is unitary, due to the fact that our Fourier transform operators can be chosen such that $\mathcal{F}^\dagger = \mathcal{F}^{-1}$, and thus

$$\begin{aligned} &\boldsymbol{U}\boldsymbol{U}^\dagger = \\ &\mathcal{F}^{-1}\left(\frac{k^2 - \alpha^\star k_0^2}{k^2 - \alpha k_0^2}\boldsymbol{\Pi}_T + \frac{\alpha^\star}{\alpha}\boldsymbol{\Pi}_L\right)\left(\frac{k^2 - \alpha k_0^2}{k^2 - \alpha^\star k_0^2}\boldsymbol{\Pi}_T + \frac{\alpha}{\alpha^\star}\boldsymbol{\Pi}_L\right)\mathcal{F} \\ &\qquad = \mathbb{1}_3, \end{aligned} \tag{S16}$$

where we used the fact that $\boldsymbol{\Pi}_T\boldsymbol{\Pi}_L = 0$, $\boldsymbol{\Pi}_T\boldsymbol{\Pi}_T = \boldsymbol{\Pi}_T$, and $\boldsymbol{\Pi}_L\boldsymbol{\Pi}_L = \boldsymbol{\Pi}_L$.

Note that the unitary operator only changes the complex argument of the transverse and longitudinal projected components. It can thus be seen from equation (S14) that the eigenvalue of $\overleftrightarrow{G}$ with the largest absolute value equals $\frac{\mathrm{i}}{\alpha_\mathrm{i} k_0^2}$, and that the real part of the eigenvalue must be no larger than $\frac{1}{2\alpha_\mathrm{i} k_0^2}$ in absolute value, while its imaginary part cannot be negative. More specifically, all eigenvalues are contained within a disk of radius $\frac{1}{2\alpha_\mathrm{i} k_0^2}$ around the point $\frac{\mathrm{i}}{2\alpha_\mathrm{i} k_0^2}$ in the complex plane.

We note that for numerical purposes it is space and time-efficient to implement the dyadic Green function operation in k-space. The resulting multiplication only requires space to store the Fourier-transformed input, output, and Green function. The vector operations of the latter can be implemented without constructing the full, multi-dimensional, matrix for the dyadic Green function, though at least a scalar array, $|\boldsymbol{k}|^2$, for normalizing the k-vectors must be stored.

S1C Computation and memory efficiency

The computational efficiency of the modified Born series iteration has been discussed previously by Osnabrugge et al.[2]. In general the convergence of the iteration is approximately inversely proportional to the range of susceptibilities in the calculation volume. The method is therefore not efficient for calculations that involve metals. The anisotropic algorithm is no different in this respect. It is most efficient for heterogeneous dielectric materials such as biological tissue. However, the anisotropic version does require the addition of 3-vectors and multiplications of 3×3 matrices instead of scalar operations. The algorithm for anisotropic permittivity can thus be expected to be 9 times slower than the scalar wave algorithm[2], and 3 times slower than the isotropic vector algorithm[3]. It should also be noted that simulating inhomogeneous magnetic properties introduces the discretized differential operator, $\mathcal{D}$. This largest singular value of this operator depends on the sampling density of the computation volume. This is equivalent to having a large variation in optical properties in the sample, which is known to lead to significantly slower convergence of the modified Born series[2].

The main limitation of any large scale electro-magnetic calculation is computer memory. The advantage of the presented algorithm is that the required memory scales with the number of sample points, P. This is important, considering that the calculation volumes of interest can be several orders of magnitude larger than the wavelength in all three dimensions. At a very minimum the electro-magnetic field in the calculation volume must be stored. The magnetic field can be calculated from the electric field, hence it is sufficient to store only the electric vector field, $\boldsymbol{E}$. This can be done using $3P$ complex floating point numbers to represent a single frequency field in the calculation volume. Each iteration calculates a correction term for the field, $\Delta\boldsymbol{E}$, thus bringing the total to $6P$.

Anisotropic permittivity can be represented by a 3×3 matrix for each point in space, to make up a block-diagonal matrix of dimension $3P \times 3P$. However, this can be efficiently stored using $9P$ complex numbers, while all matrix operations, as well as the fast-Fourier transforms can be performed in-place. It is not necessary to explicitly calculate the matrix for the dyadic Green's function (S1B), though it is necessary to determine the normalization factor $\|\boldsymbol{k}\|$ for all k-vectors. This requires storage for P values.

Anisotropy in the permeability, $\boldsymbol{\mu}$, or the coupling factors, $\boldsymbol{\xi}$ and $\boldsymbol{\zeta}$, will increase the memory requirements accordingly. If the source current $\boldsymbol{S}$ is not sparse, a further $3P$ complex values need to be stored. The memory requirements are summarized in Table S1 for the case of a spatially variant $\boldsymbol{\epsilon}$, and optionally, a spatially variant $\boldsymbol{\mu}$, $\boldsymbol{\xi}$, and $\boldsymbol{\zeta}$. The memory requirements are listed in different columns for the isotropic and the anisotropic cases. Mixed cases are omitted for brevity.

Table S1: **Storage requirements for the algorithm.**

spatial variability	isotropic	anisotropic
$\boldsymbol{\epsilon}$	$11P$	$19P$
$\boldsymbol{\epsilon}$, and $\boldsymbol{\mu}$	$12P$	$28P$
$\boldsymbol{\epsilon}$, $\boldsymbol{\mu}$, $\boldsymbol{\xi}$, and $\boldsymbol{\zeta}$	$14P$	$46P$

S1D Sampling and prevention of aliasing

The sampling grid size and density are important considerations for the calculation accuracy. In Algorithm S1, every iteration requires a multiplication by the preconditioned general susceptibility, $\boldsymbol{\chi}$, an addition of the source, $\boldsymbol{S}$, a convolution with the dyadic Green's function, and another multiplication by $\boldsymbol{\chi}$. Multiplications and additions cannot increase the spatial extent of the field $\boldsymbol{E}$ beyond the union of that of its left and right hand sides. However, the Green function has an infinite extent and therefore also the result of its convolution with a finite field. When the convolution operation is implemented as a multiplication in Fourier space, periodic boundary conditions are implicitly assumed. Alternative boundary conditions can be readily simulated by defining layers of non-transmitting material at the boundary. The sampling grid size must therefore be sufficiently large to fit both the volume of interest and multi-cell boundaries that adequately absorb the field.

To discuss the sampling density, we consider the Fourier transform of the iteration update calculation:

$$\Delta\widetilde{\boldsymbol{E}} \quad = \quad \frac{\mathrm{i}}{\alpha_\mathrm{i}} \widetilde{\boldsymbol{\chi}} * \left[\widetilde{\boldsymbol{G}} \left(\widetilde{\boldsymbol{\chi}} * \widetilde{\boldsymbol{E}} + \widetilde{\boldsymbol{S}} \right) - \widetilde{\boldsymbol{E}} \right]. \tag{S17}$$

In Fourier space, each step requires a convolution, $*$, with the Fourier transform of the susceptibility, $\widetilde{\boldsymbol{\chi}}$, followed by an addition with the Fourier transform of the source, $\widetilde{\boldsymbol{S}}$, a multiplication with the Fourier transform of the dyadic Green's function, $\widetilde{\boldsymbol{G}}$, and another convolution with $\widetilde{\boldsymbol{\chi}}$. Although the multiplication by $\widetilde{\boldsymbol{G}}$, tends to suppress high spatial frequencies, it does not strictly limit the bandwidth support of the product. On the other hand, convolutions with $\widetilde{\boldsymbol{\chi}}$ do extend the bandwidth support twice per iteration. If we define $W_{\boldsymbol{\chi}}$, $W_{\boldsymbol{E}}$, and $W_{\boldsymbol{S}} \leq W_{\boldsymbol{\chi}} + W_{\boldsymbol{E}}$ as the spatial-frequency band-widths of $\boldsymbol{\chi}$, $\boldsymbol{E}$, and $\boldsymbol{S}$, respectively; it can seen that the bandwidth of the iteration update, $\Delta\boldsymbol{E}$, must be $W_{\Delta\boldsymbol{E}} \leq 2W_{\boldsymbol{\chi}} + W_{\boldsymbol{E}}$. Therefore, even when both the material properties and the source are smooth functions with a finite bandwidth, the sampling density of the calculation must, in principle, be increased by $2W_{\boldsymbol{\chi}}$ with every iteration step.

In practice, the calculation must be performed on a sample grid with finite bandwidth, W, and sample spacing, W^{-1}. With the notable exception of superoscillations, as long as the material properties and source field are spatially band-limited, the solution for the electric field can be expected to be concentrated around the Ewald sphere. We found that the smoothing effect of the Green function is generally sufficient to suppress the highest spatial frequencies that are affected by aliasing. When we accept the approximation that the solution must be band-limited, aliasing can be completely eliminated by low-pass filtering the field after each iteration step. To achieve this, the calculation must be performed with a sampling band-width $W \geq W_{\boldsymbol{\chi}} + W_{\boldsymbol{E}}$, in other words, with a sampling density that is no smaller than the sum of the Nyquist rates for $\boldsymbol{E}$ and $\boldsymbol{\chi}$. The two convolutions with $\widetilde{\boldsymbol{\chi}}$ expand the support in frequency space to $2W_{\boldsymbol{\chi}} + W_{\boldsymbol{E}}$, thereby causing aliasing in the upper $W_{\boldsymbol{\chi}}$-part of the band. However, the lowest spatial frequencies in the $W_{\boldsymbol{E}}$-band are not affected. All aliasing artefacts can thus

be avoided by eliminating all but the spatial frequencies within the lower W_E-band of the iteration update, $\Delta \boldsymbol{E}$. Since the suppression of the highest spatial frequencies can be applied at the end of each iteration as a projection onto a subspace, it can be seen that convergence must also hold in this sub-space.

It should also be noted that the algorithm as it is presented here requires the material properties to be sampled on a regular grid. This enables efficient convolutions using the fast-Fourier transform and it simplifies the implementation in general. However, one could imagine structures that require a higher sampling density in specific regions. To address such need, we envision that the method can be extended to irregular grids by using non-uniform Fourier transforms to perform the convolutions[4,5].

S1E Robustness to errors

Unlike time-stepped methods such as FDTD, the iteration presented here is robust to numerical errors. The N^{th}-correction term, $\Delta \boldsymbol{E}_N$, is obtained from recursive equation (S1):

$$\begin{aligned} \boldsymbol{E}_N &= \boldsymbol{\Gamma} \overset{\leftrightarrow}{\boldsymbol{G}} * k_0^2 \boldsymbol{\chi} \boldsymbol{E}_{N-1} - \boldsymbol{\Gamma} \boldsymbol{E}_{N-1} + \boldsymbol{E}_{N-1} + \boldsymbol{\Gamma} \overset{\leftrightarrow}{\boldsymbol{G}} * \boldsymbol{S} \\ \Rightarrow \Delta \boldsymbol{E}_N &= \boldsymbol{E}_N - \boldsymbol{E}_{N-1} = \boldsymbol{\Gamma} \left[\overset{\leftrightarrow}{\boldsymbol{G}} * \left(k_0^2 \boldsymbol{\chi} \boldsymbol{E}_{N-1} + \boldsymbol{S} \right) - \boldsymbol{E}_{N-1} \right]. \end{aligned}$$

For the initial field estimate, If we begin with a null field estimate $\boldsymbol{E}_0 = \boldsymbol{0}$, this iteration is equivalent to the modified series (S1). However, it should be noted that when $\boldsymbol{E}_0 \neq \boldsymbol{0}$, the iteration corresponds to a different series, yet one that converges under the same conditions and to the same limit:

$$\begin{aligned} \boldsymbol{E}_N &= \boldsymbol{\Gamma} \left[\overset{\leftrightarrow}{\boldsymbol{G}} * \left(k_0^2 \boldsymbol{\chi} \boldsymbol{E}_{N-1} + \boldsymbol{S} \right) - \boldsymbol{E}_{N-1} \right] + \boldsymbol{E}_{N-1} \\ &= \boldsymbol{M} \boldsymbol{E}_{N-1} + \boldsymbol{\Gamma} \overset{\leftrightarrow}{\boldsymbol{G}} * \boldsymbol{S} \\ &= \boldsymbol{M}^N \boldsymbol{E}_0 + \left[\sum_{p=0}^{N-1} \boldsymbol{M}^p \right] \boldsymbol{\Gamma} \overset{\leftrightarrow}{\boldsymbol{G}} * \boldsymbol{S} \end{aligned} \tag{S18}$$

It can be seen that the N^{th}-term of the series differs by $\boldsymbol{M}^N \boldsymbol{E}_0$ from that of series (S1), and that its corresponding residue, $\|\boldsymbol{E} - \boldsymbol{E}_N\|$, is

$$\begin{aligned} &\left\| \left[\sum_{p=0}^{\infty} \boldsymbol{M}^p \right] \boldsymbol{\Gamma} \overset{\leftrightarrow}{\boldsymbol{G}} * \boldsymbol{S} - \boldsymbol{M}^N \boldsymbol{E}_0 - \left[\sum_{p=0}^{N-1} \boldsymbol{M}^p \right] \boldsymbol{\Gamma} \overset{\leftrightarrow}{\boldsymbol{G}} * \boldsymbol{S} \right\| \\ &\qquad = \left\| \boldsymbol{M}^N \left(\boldsymbol{E} - \boldsymbol{E}_0 \right) \right\|. \end{aligned} \tag{S19}$$

Provided that the absolute largest eigenvalue of $\boldsymbol{M}$ is smaller than one, the upper bound on the residue is tightened with every iteration, independently of the choice of the initial field. However, it can be seen from equation (S19) that starting from an approximate solution has a similar effect as skipping the first iterations, and allowing the algorithm to reach convergence in less time. The corrections in consecutive terms also prevent the accumulation of numerical errors, a common issue with techniques such as the finite-difference time-domain method.

Since the iteration does not have to be started from the all-zero field, convergence may be reached faster if an approximate solution is provided as the initial field of the iteration. It is straightforward to start the algorithm with an approximate solution to the corresponding isotropic or scalar problem. Techniques such as the beam-propagation method or two-dimensional fast marching[6], could be leveraged as an initialization of the anisotropic algorithm so to reduce the total runtime.

S2 Convergence of the modified Born series

S2A Geometrical Interpretation

The iteration as given by equation (S1) calculates a series of correction terms of the form $\boldsymbol{M}^p \boldsymbol{E}_0$. Independently of the initial value $\boldsymbol{E}_0$, this series is guaranteed to converge when all eigenvalues of $\boldsymbol{M}$ are less than one in absolute value. Recall that the preconditioned iteration operator, $\boldsymbol{M}$, is given by definition (S2).

In the previous section it is shown with equation (S14) that the dyadic Green function, $\overset{\leftrightarrow}{\boldsymbol{G}}$, can be decomposed in the scaled sum of an identity and a unitary transformation, $\boldsymbol{U}$:

$$\overset{\leftrightarrow}{\boldsymbol{G}} \equiv \frac{\mathbb{1}_3 - \boldsymbol{U}}{2} \frac{\mathrm{i}}{\alpha_\mathrm{i} k_0^2}, \tag{S20}$$

Substitution in equation (S2) allows it to written in terms of just the preconditioner, $\boldsymbol{\Gamma}$, and the susceptibility, $\boldsymbol{\chi}$:

$$\boldsymbol{M} \equiv \Gamma \left(\frac{\mathbb{1}_3 - \boldsymbol{U}}{2} \frac{\mathrm{i}}{\alpha_\mathrm{i}} \boldsymbol{\chi} - \mathbb{1}_3 \right) + \mathbb{1}_3. \tag{S21}$$

The operator $\boldsymbol{\Gamma}$ must be chosen so that the largest eigenvalue of the preconditioned iteration operator, $\boldsymbol{M}$, is less than 1 in absolute value. In other words, $\left| \boldsymbol{E}_{\max}^\dagger \boldsymbol{M} \boldsymbol{E}_{\max} \right| < 1$, where $\boldsymbol{E}_{\max}$ is the eigenvector of the largest eigenfunction of $\boldsymbol{M}$. We will now investigate the effects of these operations using a geometrical construction of the complex values that $\boldsymbol{E}_{\max}^\dagger \boldsymbol{M} \boldsymbol{E}_{\max}$ can take and show under what conditions these fall within the unit disk around the origin in the complex plane.

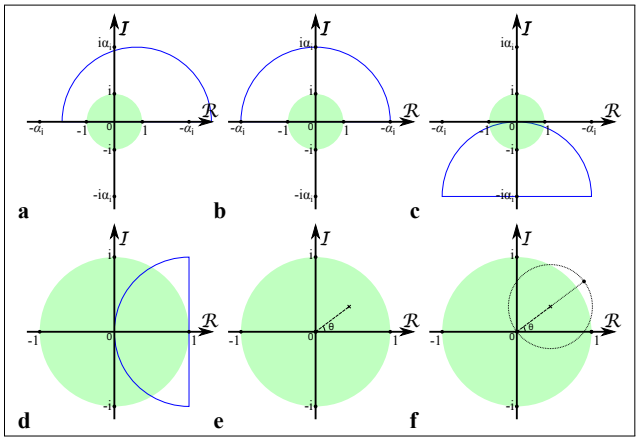

Figure S1: Geometrical interpretation of the multiplication by the susceptibility, $\boldsymbol{\chi}$, and the dyadic Green function, $\overset{\leftrightarrow}{\boldsymbol{G}}$. Values are shown in the complex plane and the final values must be within the unit disk (green) to ensure convergence. The blue semi-circle indicates the possible values of $\boldsymbol{E}_{\max}^\dagger \boldsymbol{\epsilon} \boldsymbol{E}_{\max}$ (a), $\boldsymbol{E}_{\max}^\dagger \left(\boldsymbol{\epsilon} - \alpha_\mathrm{r} \mathbb{1}_3 \right) \boldsymbol{E}_{\max}$ (b), $\boldsymbol{E}_{\max}^\dagger \boldsymbol{\chi} \boldsymbol{E}_{\max}$ (c), and , $\boldsymbol{E}_{\max}^\dagger \frac{\mathrm{i}}{\alpha_\mathrm{i}} \boldsymbol{\chi} \boldsymbol{E}_{\max}$ (d), associated with the susceptibility. (e, f) The dotted circles with dashed radius represents the decomposition of the dyadic Green function in the sum of an identity (dashed line between origin and $\times$) and unitary operation (dotted line and circle around point marked with $\times$).

The expression $\boldsymbol{E}_{\max}^\dagger \boldsymbol{M} \boldsymbol{E}_{\max}$ is constructed step-by-step, starting from the permittivity, $\boldsymbol{\epsilon} \equiv \boldsymbol{\chi} + \alpha \mathbb{1}_3$. The real and imaginary parts of $\boldsymbol{E}_{\max}^\dagger \boldsymbol{\epsilon} \boldsymbol{E}_{\max}$ can be seen to be:

$$\begin{aligned} \mathrm{Re} \left\{ \boldsymbol{E}_{\max}^\dagger \boldsymbol{\epsilon} \boldsymbol{E}_{\max} \right\} &= \boldsymbol{E}_{\max}^\dagger \frac{\boldsymbol{\epsilon} + \boldsymbol{\epsilon}^\dagger}{2} \boldsymbol{E}_{\max} \quad \text{and} \\ \mathrm{Im} \left\{ \boldsymbol{E}_{\max}^\dagger \boldsymbol{\epsilon} \boldsymbol{E}_{\max} \right\} &= \boldsymbol{E}_{\max}^\dagger \frac{\boldsymbol{\epsilon} - \boldsymbol{\epsilon}^\dagger}{2\mathrm{i}} \boldsymbol{E}_{\max}, \end{aligned} \tag{S22}$$

where gain-free media have a positive definite imaginary part. Assuming that the medium is free from gain and finite-valued, then there must exist a semi-circle as shown in blue in Figure S1(a) that contains all values $\boldsymbol{E}_{\max}^\dagger \boldsymbol{\epsilon} \boldsymbol{E}_{\max}$ in the upper half

of the complex plane. The values $\boldsymbol{E}_{\max}^{\dagger}\left(\boldsymbol{\epsilon}-\alpha_{\mathrm{r}}\mathbb{1}_3\right)\boldsymbol{E}_{\max}$ are contained within the same semi-circle, now centered at the origin as shown in Figure S1(b). Although $\boldsymbol{E}_{\max}$ is not necessarily also an eigenfunction of $\boldsymbol{\chi}$, it can be seen that when $\boldsymbol{\epsilon}$ is positive definite, such semi-circle must exist with radius, α_{i}, equal to the largest singular value of $\boldsymbol{\epsilon}-\alpha_{\mathrm{r}}\mathbb{1}_3$. The values $\boldsymbol{E}_{\max}^{\dagger}\left(\boldsymbol{\epsilon}-\alpha\mathbb{1}_3\right)\boldsymbol{E}_{\max}=\boldsymbol{E}_{\max}^{\dagger}\boldsymbol{\chi}\boldsymbol{E}_{\max}$ can be seen to be contained within the semi-circle centered at $-\mathrm{i}\alpha_{\mathrm{i}}$ and with its circumference touching the origin, as shown in Figure S1(c). The multiplication by $\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}$ normalizes and rotates the semi-circle 90° around the origin for the values $\boldsymbol{E}_{\max}^{\dagger}\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_{\max}$ as shown in Figure S1(d). The radius of the semi-circle is now 1, and its base diagonal parallel to the imaginary axis through the point $1+0\mathrm{i}$ on the real axis.

When the dyadic Green function tensor is applied to a point in Hilbert space, the result is contained within the sphere with as diagonal the line-segment between the origin and the original point, scaled by $\frac{1}{\alpha_{\mathrm{i}}k_0^2}$. The scaling factor has already been treated together with the susceptibility. To understand the effect of the dyadic Green function, we decompose $\boldsymbol{E}_{\max}$ in the eigenfunction basis of $\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}$ as $\boldsymbol{E}_{\max}=\sum_i c_i\boldsymbol{E}_i$, and first consider its effect on the eigenfunctions $\boldsymbol{E}_i$ of $\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}$ with eigenvalues λ_i. Figure S1(e) shows how the circle's center point is constructed by the identity term in the dyadic Green function $\boldsymbol{E}_i^{\dagger}\frac{\mathbb{1}_3}{2}\left(\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_i\right)$, while the inclusion of the unitary term, $\boldsymbol{E}_i^{\dagger}\frac{\mathbb{1}_3-\boldsymbol{U}}{2}\left(\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_i\right)$, outlines the circle's circumference as shown in Figure S1(f). Note that the unitary transformation may reduce the magnitude of the final projection, hence the circle only indicates the bounds of a disk of possible values.

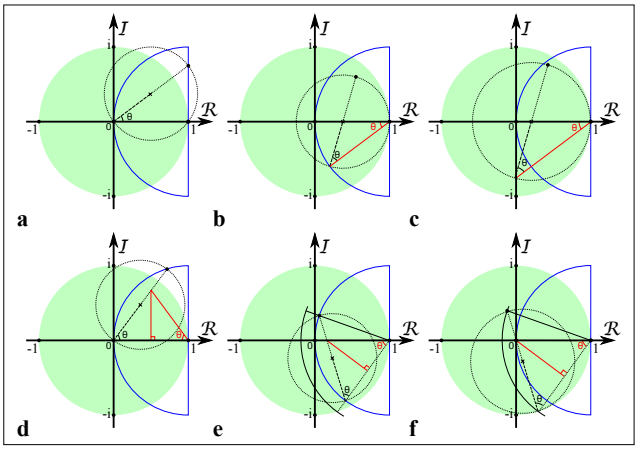

Figure S2: Geometrical interpretation of the converging iteration. The possible values that the expression $\boldsymbol{E}_i^{\dagger}\boldsymbol{M}\boldsymbol{E}_i$ can take are constructed in the complex plane. The susceptibilities can be separated into two limiting cases: purely reactive susceptibility without absorption represented by the diagonal of the semi-circle (a-c), and those corresponding to a permittivity with an absolute value α_{i} on the circumference (d-f). The red lines indicate geometrical constructions of equilateral triangles. (a,d) The values for $\boldsymbol{E}_i^{\dagger}\frac{\mathbb{1}_3-\boldsymbol{U}}{2}\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_i$ must lie within the black dotted circle with as diagonal the line segment between the origin and any point in the blue semi-circle. The preconditioner $\boldsymbol{\Gamma}$ rotates (b,e) and scales (c,f) the values around the point $1+0\mathrm{i}$ on the real axis.

The combined effects of the susceptibility and the dyadic Green function allow us to study the bounds on the eigenvalues of the preconditioned iteration operation $\boldsymbol{M}$. The value of $\boldsymbol{E}_i^{\dagger}\frac{\mathbb{1}_3-\boldsymbol{U}}{2}\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_i$ can be inside any disk between the origin and any point within the previously described semi-circle. Figure S2(a) shows the semi-circle in blue and the radii and circumference of a disk for one specific example in black dotted line. It can be noted that some values of $\boldsymbol{E}_i^{\dagger}\frac{\mathbb{1}_3-\boldsymbol{U}}{2}\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_i$ still fall outwith the unit circle in the complex plane, marked as a green background. This lies at the basis of the divergence and emphasizes the need for the preconditioner. The final three operations to construct $\boldsymbol{E}_i^{\dagger}\boldsymbol{M}\boldsymbol{E}_i$ as defined in equation (S21) are a subtraction of the identity, a multiplication by the preconditioner $\boldsymbol{\Gamma}$, and the re-addition of the identity. Since $\boldsymbol{\Gamma}$ is a linear operation, it can be understood as a scaling and rotation around the origin in the complex plane, and the three combined operations are the same scaling and rotation, but around the point $1+0\mathrm{i}$. We will now show that this ensures that all values fall within the unit disk. The black dotted circle shown in Figure S2(a) goes through the point $1+0\mathrm{i}$ when $\boldsymbol{E}_i^{\dagger}\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_i$ coincides with the base of the blue semi-circle because the base and the real axis form a right-angle triangle with the diagonal of the black dotted circle. To ensure that all points of the black dotted circle shown in Figure S2(a) are on the unit disk, one has to rotate it around the point $1+0\mathrm{i}$ and make it tangent to the unit disk. It can be seen that this can only be achieved when the operation $\boldsymbol{\Gamma}$ introduces a phase θ that is equal to that introduced by $\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}$ in the previous step. In other words, θ is the complex argument of the previously studied expression $\boldsymbol{E}_i^{\dagger}\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_i$. This operation rotates the origin of the black dotted circle onto the real axis. The multiplication by $\boldsymbol{\Gamma}$ can also scale the result. The scaling factor should be positive and sufficiently small so that the resulting radius remains below one. When the same scaling factor as $\boldsymbol{E}_i^{\dagger}\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_i$ is used, the equilateral triangle is scaled to touch the imaginary axis as shown in Figure S2(c). The base of equilateral triangle (shown in red), can be seen to trace out the values $\boldsymbol{E}_i^{\dagger}\left(\mathbb{1}_3-\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\right)\boldsymbol{E}_i$ inside the unit disk. Although a larger scaling factor would be permissible in the particular case shown in Figure S1(c), it is convenient to define the preconditioner as $\boldsymbol{\Gamma}\equiv\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}$, to ensure that all the eigenvalues are contained within the unit disk.

Similarly one can see that the values $\boldsymbol{E}_i^{\dagger}\frac{\mathbb{1}_3-\boldsymbol{U}}{2}\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_i$ on the circumference of the blue semi-circle are mapped to fall into the unit disk when using the same rotation angle, θ, and scaling factor as $\boldsymbol{E}_i^{\dagger}\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_i$. Figure S2(d) shows how a different equilateral triangle can be constructed between the diagonal of the black dotted circle and the real axis. When the rotation caused by $\boldsymbol{\Gamma}$ is equal to that in the previous step, it rotates the apex of the equilateral triangle onto the real axis as shown in Figure S2(e). The black dotted circle now falls within the bounds of the unit disk. When the preconditioner $\boldsymbol{\Gamma}$ scales the result as in the previous step, the apex translates to coincide with the origin and the diagonal of the black dotted circle goes through the origin. The points connecting the base of the equilateral triangle must therefore lie on the same circle. Since one of the endpoints lies on the circumference of the unit disk, the diagonal line of the black dotted circle must fall entirely within the unit disk.

Values that are in between the two extreme cases discussed above can be written as a weighted sum. The linearity of the operation and the convexity of the unit sphere guarantees that also non-extreme values have eigenvalues that are no greater than those encountered at the boundaries. While this shows that the $\left|\boldsymbol{E}_i^{\dagger}\boldsymbol{M}\boldsymbol{E}_i\right|\leq 1$ for all eigenfunctions, $\boldsymbol{E}_i$, of $\boldsymbol{\chi}$; the same is not necessarily true for, $\boldsymbol{E}_{\max}$, an eigenfunction of $\boldsymbol{M}$ that may be a linear combinations of the eigenfunctions $\boldsymbol{E}_i$ of $\boldsymbol{\chi}$.

The general expression $\boldsymbol{E}_{\max}^{\dagger}\boldsymbol{M}\boldsymbol{E}_{\max}$ can be written as the dif-

ference

$$\boldsymbol{E}_{\max}^{\dagger}\left[\boldsymbol{\Gamma}\left(\frac{\mathbb{1}_3-\boldsymbol{U}}{2}\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}-\mathbb{1}_3\right)+\mathbb{1}_3\right]\boldsymbol{E}_{\max}= \boldsymbol{E}_{\max}^{\dagger}\left[\boldsymbol{\Gamma}\left(\frac{1}{2}\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}-\mathbb{1}_3\right)+\mathbb{1}_3\right]\boldsymbol{E}_{\max}-\boldsymbol{E}_{\max}^{\dagger}\boldsymbol{\Gamma}\frac{\boldsymbol{U}}{2}\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_{\max}. \tag{S23}$$

In the special case that the eigenfunctions of $\boldsymbol{\chi}$ are orthonormal and using $\overleftrightarrow{\boldsymbol{G}} \equiv \frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}$, the first term can be rewritten as a weighted sum over the complex values corresponding to each eigenfunction of the system using $\boldsymbol{E}_{\max} = \sum_i c_i \boldsymbol{E}_i$ with $c_i = \boldsymbol{E}_i^{\dagger}\boldsymbol{E}_{\max}$ and $\sum_i |c_i|^2 = 1$:

$$\boldsymbol{E}_{\max}^{\dagger}\boldsymbol{M}\boldsymbol{E}_{\max}= \sum_i |c_i|^2\left[\frac{\lambda_i^2}{2}-\lambda_i+1\right]-\frac{1}{2}\left(\boldsymbol{\Gamma}^{\dagger}\boldsymbol{E}_{\max}\right)^{\dagger}\boldsymbol{U}\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_{\max}. \tag{S24}$$

Here, the summation term defines the center position of the circle for $\boldsymbol{E}_{\max}$ as a weighted average of the center positions of the circles corresponding to the eigenfunctions $\boldsymbol{E}_i$. Since for each eigenfunction, a circle with radius $\frac{1}{2}|\lambda_i|^2$ fits within the unit disk of the complex plane so the absolute value of the summation term is limited by the inequality

$$\left|\sum_i |c_i|^2\left[\frac{\lambda_i^2}{2}-\lambda_i+1\right]\right| \le \sum_i |c_i|^2\left[1-\frac{|\lambda_i|^2}{2}\right]= \sum_i |c_i|^2-\sum_i |c_i|^2\frac{|\lambda_i|^2}{2}=1-\frac{1}{2}\sum_i |c_i|^2|\lambda_i|^2. \tag{S25}$$

This leaves sufficient space to fit a circle of radius $\frac{1}{2}\sum_i |c_i|^2|\lambda_i|^2$ within the unit disk of the complex plane. The second term in equation (S24) is one-half of the dot product of a unitary operation and two terms with identical l_2-norm $\left\|\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}\boldsymbol{E}_{\max}\right\| = \sqrt{\sum_i |c_i|^2|\lambda_i|^2} = \left\|\boldsymbol{\Gamma}^{\dagger}\boldsymbol{E}_{\max}\right\|$, when $\boldsymbol{\Gamma}$ is defined as $\frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}$. Hence, by defining $\boldsymbol{\Gamma} \equiv \frac{\mathrm{i}}{\alpha_{\mathrm{i}}}\boldsymbol{\chi}$, it can be seen that equation (S24) cannot be larger than 1, where α_{i} is larger than the largest singular value of $\|\boldsymbol{\Delta}\|$. The series must thus converge when the eigenfunctions of the susceptibility distribution, $\boldsymbol{\chi}$, are orthogonal, a very common situation. Yet, its eigenfunctions will not be orthogonal when the reactive, $\boldsymbol{\Delta}_{\mathrm{r}}$, and dissipative parts, $\boldsymbol{\Delta}_{\mathrm{i}}$, of $\boldsymbol{\Delta} = \boldsymbol{\chi} + \mathrm{i}\alpha_{\mathrm{i}}\mathbb{1}_3 = \boldsymbol{\Delta}_{\mathrm{r}} + \mathrm{i}\boldsymbol{\Delta}_{\mathrm{i}}$ do not commute. This would occur when a birefringent crystal also has a polarization dependent absorption, yet with a different axis. In what follows, such more general susceptibilities are analyzed.

S2B Numerical demonstration of the convergence

In the main text we showed that the following choice of α_{i} guarantees convergence of the modified Born series

$$\alpha_{\mathrm{i}} > \max\{A_1, A_2\}, \tag{S26}$$

where A_1 and A_2 are defined as:

$$\begin{aligned} A_1 &= \sqrt{\max_{\boldsymbol{n}}\left(\frac{1}{2}\langle\boldsymbol{n}|\,\boldsymbol{\Delta}\boldsymbol{\Delta}^{\dagger}+\boldsymbol{\Delta}^{\dagger}\boldsymbol{\Delta}\,|\boldsymbol{n}\rangle\right)} \\ A_2 &= \max_{\boldsymbol{n}}\frac{|\langle\boldsymbol{n}|\boldsymbol{\Delta}\boldsymbol{\Delta}_{\mathrm{i}}+\boldsymbol{\Delta}_{\mathrm{i}}\boldsymbol{\Delta}|\boldsymbol{n}\rangle|}{2\langle\boldsymbol{n}|\boldsymbol{\Delta}_{\mathrm{i}}|\boldsymbol{n}\rangle} \end{aligned} \tag{S27}$$

It the main text it was noted that without constraining $\boldsymbol{\Delta}_{\mathrm{i}}$ to be positive definite, A_2 can be arbitrarily large, when $\langle\boldsymbol{n}|\boldsymbol{\Delta}_{\mathrm{i}}|\boldsymbol{n}\rangle$ is close to zero. We begin this Supplementary Section by showing that this divergence can be avoided so long as $\boldsymbol{\Delta}_{\mathrm{i}}$ has an empty kernel (no eigenvectors with zero eigenvalue). Suppose that $\boldsymbol{\Delta}_{\mathrm{i}}$

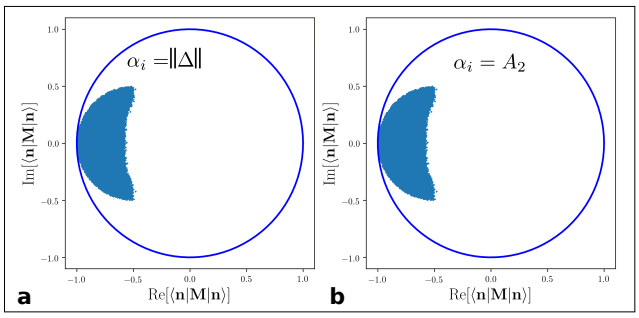

Figure S3: Numerical check that condition $\alpha_{\mathrm{i}} > \max\{A_1, A_2\}$ ensures that the numerical radius of $\boldsymbol{M}$ (given by equation (S2)) is less than unity. To produce this figure we used the condition on the numerical radius, $|\langle\boldsymbol{n}|\boldsymbol{M}|\boldsymbol{n}\rangle| < \langle\boldsymbol{n}|\boldsymbol{n}\rangle$, with 40×40 matrices. We generated 1×10^6 random matrices $\boldsymbol{\Delta}$ (positive definite $\boldsymbol{\Delta}_{\mathrm{i}}$) and unitary matrices $\boldsymbol{U}$, along with the corresponding values of α_{i} calculated as indicated in each panel (using Python's numpy library for random matrix generation[7]). For each matrix we calculated values for $|\langle\boldsymbol{n}|\boldsymbol{M}|\boldsymbol{n}\rangle|$ for a set of 40 random but orthogonal complex vectors, $\boldsymbol{n}$. The largest magnitude of these values is one of the 10^6 points plotted in each panel. In (b) (largest magnitude 0.999997) we used the condition (S32), which we know analytically to be sufficient to move the eigenvalues of $\boldsymbol{M}$ within the unit circle. In (a) (largest magnitude 0.999999) we show that $\alpha_i > \|\Delta\|$ appears to guarantee also that the eigenvalues are within the unit circle

has an eigenvector $|\boldsymbol{n}_0\rangle$ with eigenvalue λ, and consider $|\boldsymbol{n}\rangle = |\boldsymbol{n}_0\rangle + \eta|\boldsymbol{n}_\perp\rangle$ where $\eta \ll 1$. The expression for A_2 inside the maximization is then given by

$$\frac{|\langle\boldsymbol{n}|\boldsymbol{\Delta}\boldsymbol{\Delta}_{\mathrm{i}}+\boldsymbol{\Delta}_{\mathrm{i}}\boldsymbol{\Delta}|\boldsymbol{n}\rangle|}{2\langle\boldsymbol{n}|\boldsymbol{\Delta}_{\mathrm{i}}|\boldsymbol{n}\rangle} = \frac{\left|\begin{array}{l} 2\lambda\langle\boldsymbol{n}_0|\boldsymbol{\Delta}|\boldsymbol{n}_0\rangle+\eta\left(\langle\boldsymbol{n}_\perp|\boldsymbol{\Delta}_{\mathrm{i}}\boldsymbol{\Delta}|\boldsymbol{n}_0\rangle+\langle\boldsymbol{n}_0|\boldsymbol{\Delta}\boldsymbol{\Delta}_{\mathrm{i}}|\boldsymbol{n}_\perp\rangle\right)+ \\ +\eta\lambda\left(\langle\boldsymbol{n}_\perp|\boldsymbol{\Delta}|\boldsymbol{n}_0\rangle+\langle\boldsymbol{n}_0|\boldsymbol{\Delta}|\boldsymbol{n}_\perp\rangle\right)+\eta^2\langle\boldsymbol{n}_\perp|\boldsymbol{\Delta}\boldsymbol{\Delta}_{\mathrm{i}}+\boldsymbol{\Delta}_{\mathrm{i}}\boldsymbol{\Delta}|\boldsymbol{n}_\perp\rangle| \end{array}\right|}{2\left(\lambda+\eta^2\langle\boldsymbol{n}_\perp|\boldsymbol{\Delta}_{\mathrm{i}}|\boldsymbol{n}_\perp\rangle\right)} \tag{S28}$$

If $\lambda = 0$ (and the kernel of $\boldsymbol{\Delta}_{\mathrm{i}}$ thus contains $|\boldsymbol{n}_0\rangle$ then the above quantity diverges as $1/\eta$ as we take η to zero. The quantity A_2 thus becomes infinite. However, if λ is non-zero (but arbitrarily small) then as $\eta \to 0$, (S28) tends to $|\langle\boldsymbol{n}_0|\boldsymbol{\Delta}|\boldsymbol{n}_0\rangle|$ which is finite, and we note, smaller than or equal to the largest singular value of $\boldsymbol{\Delta}$.

Assuming that $\langle\boldsymbol{n}|\boldsymbol{\Delta}_{\mathrm{i}}|\boldsymbol{n}\rangle$ is never zero, the quantity A_2 can be rewritten in a form that is easier to compute. First we write the positive definite Hermitian matrix $\boldsymbol{\Delta}_{\mathrm{i}}$ as the square of another Hermitian matrix $\boldsymbol{a}$

$$\boldsymbol{\Delta}_{\mathrm{i}} = \boldsymbol{a}^2 \tag{S29}$$

Defining $|\boldsymbol{n}'\rangle = \boldsymbol{a}|\boldsymbol{n}\rangle$, the quantity A_2 can be written as

$$A_2 = \max_{\boldsymbol{n}'}\frac{|\langle\boldsymbol{n}'|\boldsymbol{a}^{-1}\boldsymbol{\Delta}\boldsymbol{a}+\boldsymbol{a}\boldsymbol{\Delta}\boldsymbol{a}^{-1}|\boldsymbol{n}'\rangle|}{2\langle\boldsymbol{n}'|\boldsymbol{n}'\rangle} \tag{S30}$$

which is simply the numerical radius $r(\boldsymbol{m})$ of the matrix

$$\boldsymbol{m} = \frac{1}{2}(\boldsymbol{a}^{-1}\boldsymbol{\Delta}\boldsymbol{a}+\boldsymbol{a}\boldsymbol{\Delta}\boldsymbol{a}^{-1}). \tag{S31}$$

The numerical radius of a matrix is always less than or equal to its norm[8] $r(\boldsymbol{m}) \le \|\boldsymbol{m}\|$. In addition, the norm of a sum is never larger than the sum of the norms, and we can therefore estimate A_2 as

$$A_2 = \frac{1}{2}\left[\|\boldsymbol{a}^{-1}\boldsymbol{\Delta}\boldsymbol{a}\|+\|\boldsymbol{a}\boldsymbol{\Delta}\boldsymbol{a}^{-1}\|\right] \tag{S32}$$

where $\|\cdot\|$ indicates the l_2-norm. While the spectra of $\boldsymbol{\Delta}$ and e.g. $\boldsymbol{a}\boldsymbol{\Delta}\boldsymbol{a}^{-1}$ are the same, their norms are not. Given that the spectral radius of an operator is always less than or equal to its norm, the lowest possible value of A_2 is the magnitude of the largest eigenvalue $|\lambda_{\max}|$ of $\boldsymbol{\Delta}$. Meanwhile, A_1, being the square root of the numerical radius of a Hermitian operator, is equal to the square root of the operator norm $(1/2)\|\boldsymbol{\Delta}\boldsymbol{\Delta}^\dagger + \boldsymbol{\Delta}^\dagger\boldsymbol{\Delta}\|$, which is bounded by

$$A_1 \leq \frac{1}{2}\|\boldsymbol{\Delta}\boldsymbol{\Delta}^\dagger\| + \frac{1}{2}\|\boldsymbol{\Delta}^\dagger\boldsymbol{\Delta}\| = |\lambda_{\max}| \leq A_2 \tag{S33}$$

where the estimate for A_2 is here given by equation (S32). We can therefore use this estimate of A_2 to find an upper bound for the value of α_i. Figure S3b shows a numerical test, where we chose A_2 according to (S32) and repeatedly evaluated the inner product $\langle \boldsymbol{n}|\boldsymbol{M}|\boldsymbol{n}\rangle$ for different choices of random complex 40×40 matrices $\boldsymbol{\Delta}$ and $\boldsymbol{U}$. These tests represent $10^6 \times 40 = 4 \times 10^7$ evaluations of the inner product and appear to indicate that $\alpha_i > \|\boldsymbol{\Delta}\|$ is a tighter bound (Fig. S3a).

S3 Impedance matching and Bianisotropy

By eliminating back reflections, impedance matching ensures efficient energy transfer and communication through waveguides. Supplementary Figure S4(a-d) introduces the concept of impedance matching in one dimension. Supplementary Figures S4(a,b) shows two samples with isotropic permittivity, ϵ, and permeability, μ. Both samples contain objects (at $10\,\mu\text{m} < z < 20\,\mu\text{m}$) with identical permittivity, larger than the surrounding medium. In the left-hand panel the permeability equals the background, while in the right-hand panel it equals the permittivity for all z. Supplementary Figure S4(b) thus represents an object that is impedance matched with the surrounding medium. As can been seen from Supplementary Figure S4(c), reflections from the front and back surface interfere with the incoming wave. This is most clearly visible in the 'beating' in the intensity where the incoming and the reflected wave interfere. Supplementary Figure S4(d) shows how impedance matching successfully suppresses back reflections at both interfaces. The absence of oscillations in the field amplitude indicates an absence of back reflections.

Impedance matching also has important practical applications for simulations of infinite volumes in a finite space. Supplementary Figure S4(e) shows the electric dipole field in a box with regular absorbing boundaries and homogeneous permeability. As in the one-dimensional case, plotted in Supplementary Figure S4(a,c), significant reflections can be noted from the boundaries. Supplementary Figure S4(f) shows a dipole in a box with impedance matched absorbing boundaries. It can be seen that the interfering reflections are suppressed. In higher dimensions, impedance matching is only an approximation to perfectly matching layers. This may explain the weak, low-frequency, beating that can be noticed in the top-right corner.

The ability to account for the coupling factors enables the calculation of electromagnetic waves in materials with chiral properties. Supplementary Figures S4(g,h) show how a chiral substance slowly rotates the linear polarization of a wave that propagates through it. Supplementary Figure S4(g) shows the intensity in the horizontal and vertical polarization component, while Supplementary Figure S4(h) shows how the polarization angle changes linearly over the simulation volume's width of 10 mm.

Supplementary References

[1] B. Hecht L. Novotny. *Principles of Nano-Optics*. Cambridge University Press, June 2012.

[2] Gerwin Osnabrugge, Saroch Leedumrongwatthanakun, and Ivo M Vellekoop. A convergent born series for solving the inhomogeneous helmholtz equation in arbitrarily large media. *Journal of computational physics*, 322:113–124, 2016.

[3] Benjamin Krüger, Thomas Brenner, and Alwin Kienle. Solution of the inhomogeneous maxwell's equations using a born series. *Opt. Express*, 25(21):25165–25182, October 2017.

[4] Alok Dutt and Vladimir Rokhlin. Fast fourier transforms for nonequispaced data. *SIAM Journal on Scientific computing*, 14(6):1368–1393, 1993.

[5] Jeffrey A Fessler and Bradley P Sutton. Nonuniform fast fourier transforms using min-max interpolation. *IEEE Transactions on Signal Processing*, 51(2):560–574, 2003.

[6] Amedeo Capozzoli, Claudio Curcio, Angelo Liseno, and Salvatore Savarese. Two-dimensional fast marching for geometrical optics. *Optics express*, 22(22):26680–26695, 2014.

[7] T. E. Oliphant. *A guide to NumPy*. Trelgol Publishing, 2006.

[8] S. S. Dragomir. *Inequalities for the Numerical Radius of Linear Operators in Hilbert Spaces*. Springer, 2013.

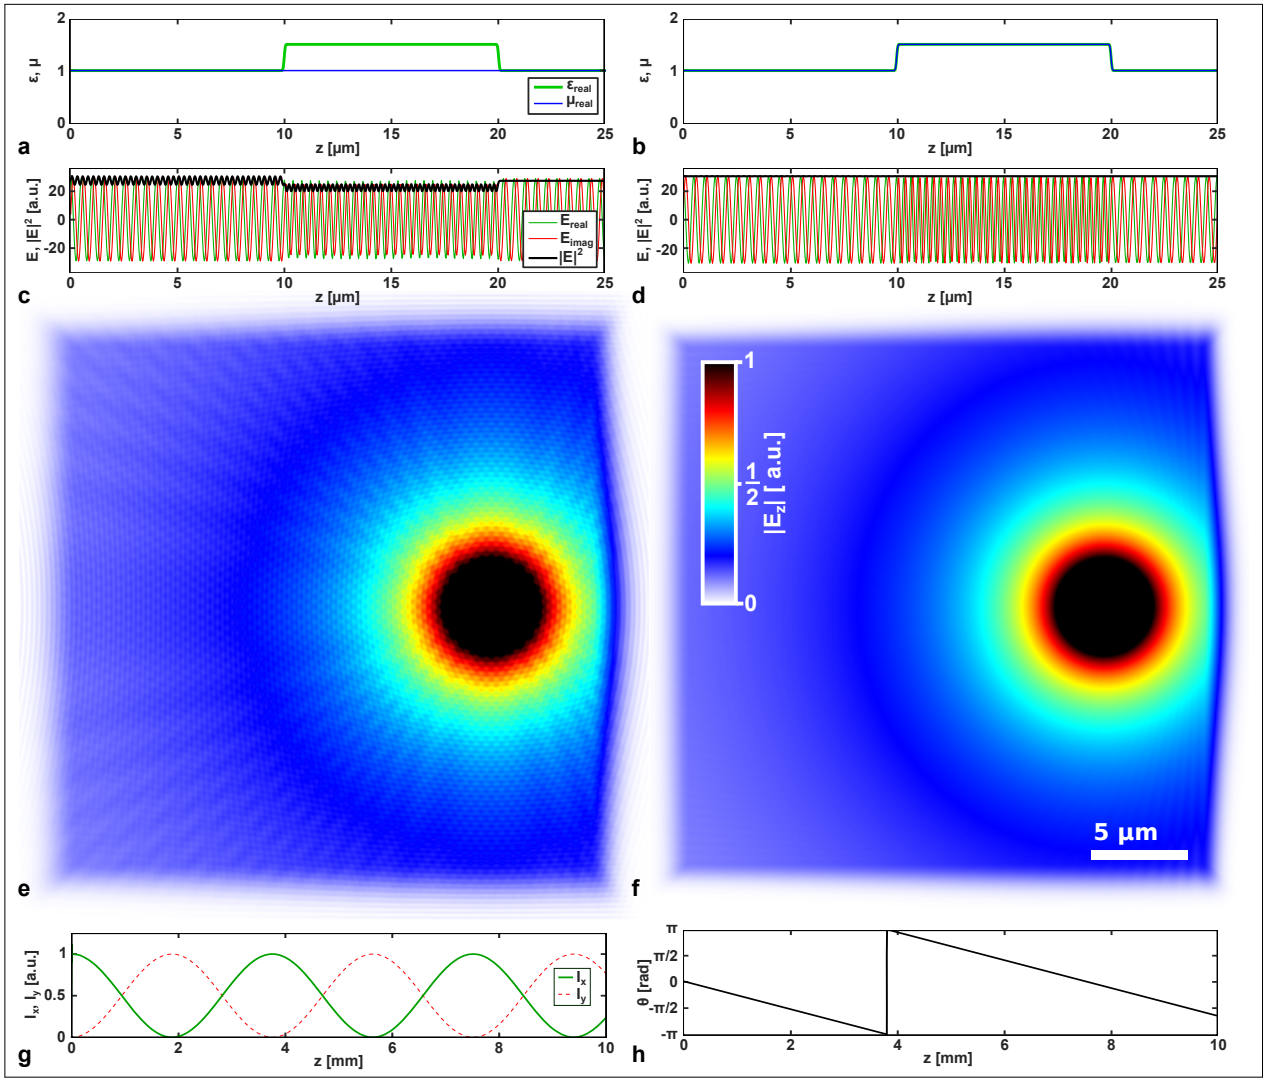

Figure S4: Demonstration of impedance matching (a-f), and propagation in a chiral medium (g,h). A plane wave in free space ($\epsilon = \mu = 1$) with a wavelength $\lambda = 500$ nm enters at $x = 10\,\mu$m from the left into dielectric slabs of thickness 10 μm, with $\mu = 1$ (a,c) and $\mu = 1.5$ (b,d). In both cases the permittivity is 1.5 (green line, panels a and b). The interference between the incoming and reflected wave is clearly visible as oscillations in intensity ($|E|^2$, black line, c). It can be seen that a fraction of the wave is reflected from the slab without impedance matching ($\mu \not\propto \epsilon$ in panel a). In contrast, a constant intensity is seen in panel (d), indicative of the absence of back-reflection for the impedance matched slab ($\mu \propto \epsilon$ in panel b). To facilitate comparison, both the intensity and field are normalized to their respective maximum value in panels (c) and (d). Panels (e,f) show the (truncated) electric field amplitude for a dipole with absorbing (e) and impedance matched (f) boundary layers. The interference with the back reflected wave, visible as beating in panel (e), is suppressed bt the impedance matched layers as seen in panel (f). (g,h) Linear polarization rotates upon propagation in a chiral medium with a high chirality that is 100 times of that of saturated glucose ($n = 1.45$, specific rotation $[\alpha]^T_{500\,\mathrm{nm}} 52.7^\circ \mathrm{mL\ g^{-1} dm^{-1}}$, at 909 g/L). The constitutive relations can thus be seen to be $\epsilon = \sqrt{1.45}, \mu = 1$, and $\xi = \zeta = 52.7\frac{909\lambda}{360}\mathrm{i} = 66.53 \times 10^{-6}\mathrm{i}$. The intensity transfer between the x-polarization (solid green) and y-polarization (dashed red) can be seen to occur several times over a propagation distance of 10 mm (g). The local angle, θ, of the linear polarization is shown in panel (h). Note the significantly larger length scale for panels (g) and (h).