

OPEN

# Data-driven acceleration of photonic simulations

Rahul Trivedi<sup>1,2\*</sup>, Logan Su<sup>1</sup>, Jesse Lu<sup>2</sup>, Martin F. Schubert<sup>2</sup> & Jelena Vuckovic<sup>1</sup> 

Designing modern photonic devices often involves traversing a large parameter space via an optimization procedure, gradient based or otherwise, and typically results in the designer performing electromagnetic simulations of a large number of correlated devices. In this paper, we investigate the possibility of accelerating electromagnetic simulations using the data collected from such correlated simulations. In particular, we present an approach to accelerate the Generalized Minimal Residual (GMRES) algorithm for the solution of frequency-domain Maxwell's equations using two machine learning models (principal component analysis and a convolutional neural network). These data-driven models are trained to predict a subspace within which the solution of the frequency-domain Maxwell's equations approximately lies. This subspace is then used for augmenting the Krylov subspace generated during the GMRES iterations, thus effectively reducing the size of the Krylov subspace and hence the number of iterations needed for solving Maxwell's equations. By training the proposed models on a dataset of wavelength-splitting gratings, we show an order of magnitude reduction (~10–50) in the number of GMRES iterations required for solving frequency-domain Maxwell's equations.

Numerical algorithms for solving Maxwell's equations are often required in a large number of design problems, ranging from integrated photonic devices to RF antennas and filters. Recent advances in large-scale computational capabilities have opened up the door for algorithmic design of electromagnetic devices for a number of applications<sup>1–3</sup>. These design algorithms typically explore very high dimensional parameter spaces using gradient-based optimization methods. A single device optimization requires ~500–1000 electromagnetic simulations, making them the primary computational bottleneck. However, during such a design process, the electromagnetic simulations being performed are on correlated permittivity distributions (e.g. permittivity distributions generated at different steps of a gradient-based design algorithm). The availability of such data opens up the possibility of using data-driven approaches for accelerating electromagnetic simulations.

Accelerating simulations in a gradient-based optimization algorithm has been previously investigated in structural engineering designs via Krylov subspace recycling<sup>4,5</sup> — the key idea in such approaches is to use the subspace spanned by simulations in an optimization trajectory to augment future simulations being performed in the same trajectory. However, these approaches become computationally infeasible if an attempt is made to exploit the full extent of the available simulation data (e.g. use simulations performed in a large number of correlated optimization trajectories) since the augmenting subspace becomes increasingly high-dimensional. Approaches that intelligently select a suitable low-dimensional subspace using machine learning models, like the one explored in this paper, are thus expected to be more efficient at exploiting the available simulation data in practical design settings. Data-driven methods for solving partial-differential equations have only recently been investigated, with demonstration of learning an 'optimal' finite-difference stencil<sup>6</sup> for time-domain simulations, or using neural networks for solving partial differential equations<sup>7</sup>.

In this paper, we investigate the possibility of accelerating finite difference frequency domain (FDFD) simulation of Maxwell's equations using data-driven models. Performing an FDFD simulation is equivalent to solving a large sparse system of linear equations, which is typically done using an iterative solver such as Generalized Minimal Residual (GMRES) algorithm<sup>8</sup>. Here we develop an accelerated solver (data-driven GMRES) by interfacing a machine learning model with GMRES. The machine learning model is trained to predict a subspace that approximates the simulation result, and the subspace is used to augment the GMRES iterations. Since the simulation is still being performed with an iterative solver, it is guaranteed that the result of the simulation will be accurate — the performance of the machine learning model only affects how fast the solution is obtained. This is a major advantage of this approach over other data-driven attempts for solving Maxwell's equations<sup>9,10</sup> that have been presented so far, in which case a misprediction by the model would result in an inaccurate simulation. Using

<sup>1</sup>E. L. Ginzton Laboratory, Stanford University, Stanford, CA, 94305, USA. <sup>2</sup>X, Mountain View, CA, 94043, USA. \*email: [rtrivedi@stanford.edu](mailto:rtrivedi@stanford.edu)

wavelength-splitting gratings as an example, we show an order of magnitude reduction in the number of GMRES iterations required for solving frequency-domain Maxwell's equations.

This paper is organized as follows - Section 1 outlines the data-driven GMRES algorithm and Section 2 presents results of applying the data-driven GMRES algorithm using two machine learning models, principal component analysis and convolutional neural networks, to simulate wavelength splitting gratings. We show that data-driven GMRES not only achieves an order of magnitude speedup against GMRES, but also outperforms a number of commonly used data-free preconditioning techniques.

## Data-driven GMRES

In the frequency domain, Maxwell's equations can be reduced to a partial differential equation relating the electric field  $\mathbf{E}(\mathbf{x})$  to its source  $\mathbf{J}(\mathbf{x})$ :

$$\left[ \nabla \times \nabla \times - \frac{\omega^2}{c^2} \varepsilon(\mathbf{x}) \right] \mathbf{E}(\mathbf{x}) = -i\omega\mu_0 \mathbf{J}(\mathbf{x}) \quad (1)$$

where  $\omega$  is the frequency of the simulation, and  $\varepsilon(\mathbf{x})$  is the permittivity distribution as a function of space. The finite difference frequency domain method<sup>11</sup> is a popular approach for numerically solving this partial differential equation – it discretizes this equation on the Yee grid with perfectly matched layers together and periodic boundary conditions used for terminating the simulation domain to obtain a system of linear equations,  $Af = b$ , where  $A$  is a sparse matrix describing the operator  $\nabla \times \nabla \times - \omega^2 \varepsilon(\mathbf{x})/c^2$  and  $b$  is a vector describing the source term  $-i\omega\mu_0 \mathbf{J}(\mathbf{x})$  on the Yee grid.

For large-scale problems, this system of equation is typically solved via a Krylov subspace-based iterative method<sup>8</sup>. Krylov subspace methods have an advantage that they only access the matrix  $A$  via matrix-vector products, which can be performed very efficiently since  $A$  is a sparse matrix. The iterative algorithm that we focus on in this paper is GMRES<sup>8</sup>. In  $i^{\text{th}}$  iteration of standard GMRES, the solution to  $Af = b$  is approximated by  $f_i$ , where:

$$f_i = \underset{f \in \mathcal{K}_i(A, b)}{\operatorname{argmin}} \|Af - b\|^2 \quad (2)$$

where  $\mathcal{K}_i(A, b) = \operatorname{span}(b, Ab, A^2b \dots A^{i-1}b)$  is the Krylov subspace of dimension  $i$  generated by  $(A, b)$ . The Krylov subspace can be generated iteratively i.e. if an orthonormal basis for  $\mathcal{K}_i(A, b)$  has been computed, then an orthonormal basis for  $\mathcal{K}_{i+1}(A, b)$  can be efficiently computed with one additional matrix-vector product. In practical simulation settings, the GMRES iteration is performed till the Krylov subspace is large enough for the residual  $r_i = \|Af_i - b\|/\|b\|$  to be smaller than a user-defined threshold. We also note that GMRES is a completely data-free algorithm – it only requires knowledge of the source vector  $b$  and the ability to multiply the matrix  $A$  with an arbitrary vector.

The number of GMRES iterations can be significantly reduced if an estimate of the solution  $A^{-1}b$  is known. To this end, we train a data-driven model on simulations of correlated structures to predict a (low-dimensional) subspace  $\mathcal{V}$  within which  $A^{-1}b$  is expected to lie. More specifically, the model predicts  $N$  vectors,  $v_1, v_2 \dots v_N$ , such that  $\mathcal{V} = \operatorname{span}(v_1, v_2 \dots v_N)$ . These vectors can then be used to augment the GMRES iterations (we refer to the augmented version of GMRES as data-driven GMRES throughout the paper) – the  $i^{\text{th}}$  iteration of data-driven GMRES can be formulated as:

$$f_i = \underset{f \in \mathcal{V} \oplus \mathcal{K}_i(\tilde{A}, \tilde{b})}{\operatorname{argmin}} \|Af - b\|^2 \quad (3)$$

where  $\tilde{A}$ , and  $\tilde{b}$  are given by:

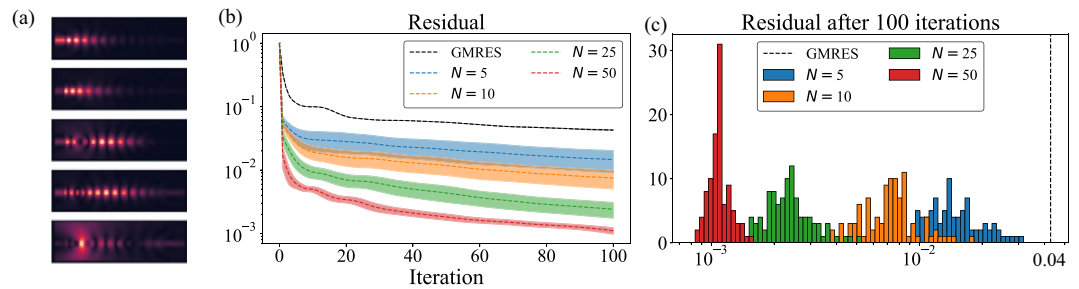
$$\tilde{A} = P_{\perp}(Av_1, Av_2 \dots Av_N)A \quad (4)$$

$$\tilde{b} = P_{\perp}(Av_1, Av_2 \dots Av_N)b \quad (5)$$

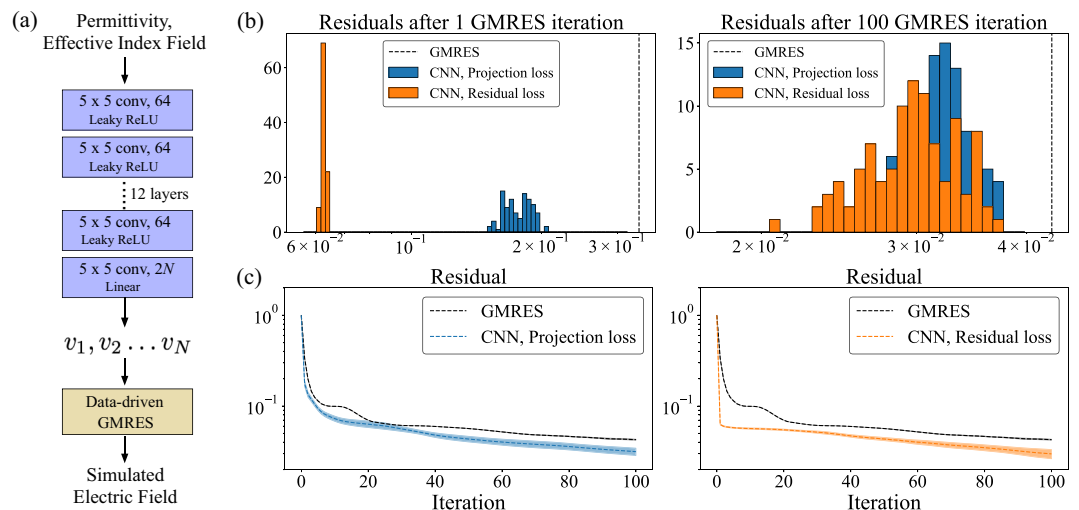
where  $P_{\perp}(Av_1, Av_2 \dots Av_N)$  is the operator projecting a vector out of the space spanned by  $Av_1, Av_2 \dots Av_N$ . Note that while in GMRES (Eq. 2) the generated Krylov subspace is responsible for estimating the entire solution  $A^{-1}b$ , in data-driven GMRES (Eq. 3) the generated Krylov subspace is only responsible for estimating the projection of  $A^{-1}b$  perpendicular to the subspace  $\mathcal{V}$ . Therefore, a large speedup in the solution of  $Af = b$  can be expected if  $\mathcal{V}$  is a good estimate of a subspace within which  $A^{-1}b$  lies. Moreover, as is shown in the supplement, an efficient update algorithm for data-driven GMRES can be formulated in a manner identical to that formulated for Generalized Conjugate Residual with inner Orthogonalization and outer Truncation (GCROT)<sup>5,12</sup>.

## Results

We investigate two data-driven models to predict the vectors  $v_1, v_2 \dots v_N$ : principal component analysis and a convolutional neural network. The dataset that we use for training and evaluating these models comprises of a collection of 2D grating splitters [Fig. 1(a)] which reflect an incident waveguide mode at  $\lambda = 1.4 \mu\text{m}$  and transmit an incident waveguide mode at  $\lambda = 1.55 \mu\text{m}$ . Throughout this paper, we focus on accelerating simulations of the grating splitters at  $\lambda = 1.4 \mu\text{m}$  – so as to train our data-driven models, we provide the dataset with the full simulations of the electric fields in the grating splitter at  $1.4 \mu\text{m}$  [Fig. 1(b)]. Additionally, we intuitively expect a well designed data-driven model to perform better if supplied with an approximation to the simulated field as



**Figure 1.** (a) Schematic of the grating splitter device that comprises the dataset. All the gratings in the dataset are  $3 \mu\text{m}$  long and are designed for a  $220 \text{ nm}$  silicon-on-insulator (SOI) platform with oxide cladding. We use a uniform spatial discretization of  $20 \text{ nm}$  while representing Eq. 1 as a system of linear equations. The resulting system of linear equations has  $229 \times 90 = 20,610$  unknown complex numbers. (b) Visualizing samples from the dataset — shown are permittivity distribution, simulated electric fields and effective index fields for 4 randomly chosen samples. All fields are shown at a wavelength of  $1.4 \mu\text{m}$ .

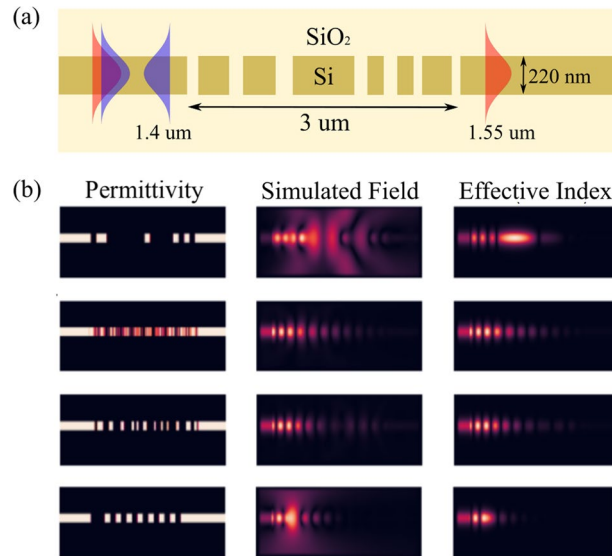


**Figure 2.** (a) First five principal components of the electric fields in the grating splitter dataset. (b) Performance of data-driven GMRES on the evaluation dataset when supplied with different number of principal components ( $\sim 200$  samples from the training set were used for computing the principal components) — the dotted line shows the mean residual, and the solid colored background indicates the region within one standard deviation around the mean residual. (c) Histogram of the residual after 100 data-driven GMRES iterations for different  $N$  computed over 100 randomly chosen samples from the evaluation dataset. The black vertical dashed line indicates the mean residual after 100 iterations of GMRES over the evaluation dataset.

input for predicting the subspace  $\mathcal{V}$ . To this end, we provide the dataset with effective index simulations<sup>13</sup> of the electric fields in the grating splitter [Fig. 1(b)]. The effective index simulations are very cheap to perform since they are equivalent to solving Maxwell's equations in 1D, making them an attractive approximation to the simulated field that the data-driven model can exploit.

**Principal component analysis.** The first data-driven model that we consider for accelerating FDFD simulations is using the principal components<sup>14</sup> computed from the simulated fields in the training dataset as  $v_1, v_2 \dots v_N$ . The first 5 principal components of the training dataset are shown in Fig. 2(a) — we computed these principal components by performing an incomplete singular value decomposition of a matrix formed with the electric field vectors of the training dataset as its columns. The first two principal components appear like fields “reflected” from the grating region, whereas the higher order principal components capture fields that are either transmitted or scattered away from the grating devices. Note that the principal components are not necessarily solutions to Maxwell's equations for a grating structure, but provide an estimate of a basis on which the solutions of Maxwell's equations for the grating structures can be accurately represented.

Using principal components as  $v_1, v_2 \dots v_N$  in data-driven GMRES, Fig. 2(b) shows the residual  $r_i = \|Af_i - b\|/\|b\|$  as a function of the number of iterations  $i$  and Fig. 2(c) shows the histogram of the residual after 100 iterations over the evaluation dataset. We clearly see an order of magnitude speed up in convergence rate for GMRES when supplemented with  $\geq 5$  principal components. Note that a typical trajectory of data-driven GMRES shows a significant reduction in the residual in the first iteration. This corresponds to GMRES finding the most suitable vector



**Figure 3.** (a) Schematic of the CNN based data-driven GMRES — a convolutional neural network takes as input the permittivity and effective index field and produces as an output the vectors  $v_1, v_2 \dots v_N$ . These vectors are then supplied to the data-driven GMRES algorithm, which produces the full simulated field. (b) Histogram of the residual after 1 and 100 data-driven GMRES iterations evaluated over the evaluation dataset. We consider neural networks trained with both the projection loss function  $l_{\text{proj}}$  and residual loss function  $l_{\text{res}}$ . The vertical dashed lines indicate the mean residual after 1 and 100 iterations of GMRES over the evaluation dataset. (c) Performance of the data-driven GMRES on the evaluation dataset when supplied with the vectors at the output of the convolutional neural networks trained with the projection loss function  $l_{\text{proj}}$  and the residual loss function  $l_{\text{res}}$ . The dotted line shows the mean residual, and the solid colored background indicates the region within  $\pm$  standard deviation around the mean residual.

minimizing the residual within the space spanned by  $b$  and the supplied principal components. Moreover, the residuals in the data-driven GMRES decrease more rapidly than in GMRES. This acceleration can be attributed to the fact that the Krylov subspace generated corresponds to the matrix  $\tilde{A}$  defined in Eq. 4 instead of  $A$ . Since  $A$  corresponds to a double derivative operator, its application is equivalent to convolving the electric field on the grid with a  $3 \times 3$  filter. Therefore,  $A$  generates a Krylov subspace almost on a pixel by pixel basis<sup>15</sup>. On the other hand, application of  $\tilde{A}$  to a vector is equivalent to application of  $A$  (which is a  $3 \times 3$  filter) followed by the projection operator  $P_{\perp}(Av_1, Av_2 \dots Av_N)$  (which is a fully dense operator).  $\tilde{A}$  can thus generate a Krylov subspace that spans the entire simulation region within the first few iterations, leading to a larger decrease in the residual  $r_i$  in data-driven GMRES as compared to GMRES.

We also reemphasize the fact that the computation of principal components needs to be performed only once for a given training dataset. Additionally, we observed that we did not need to use the entire dataset (which had  $\sim 22,500$  electric field vectors) for computing the principal components — using  $\sim 200$  randomly chosen data samples already provided a good estimate of the dominant principal components. Consequently, the computational cost of data-driven GMRES is still dominated by the iterative solve which indicates that the residual obtained by data-driven GMRES when compared to GMRES is a good measure of the obtained speedup.

**Convolutional neural network.** While using principal components achieves an order of magnitude speed up over GMRES, this approach predicts the same subspace  $\mathcal{V}$  irrespective of the permittivity distribution being simulated. Intuitively, it might be expected that a data-driven model that specializes  $\mathcal{V}$  to the permittivity distribution being simulated would unlock an even greater speed up over GMRES. To this end, we train a convolutional neural network that takes as input the permittivity distribution of the grating device as well as the effective index electric field and predicts the vectors  $v_1, v_2 \dots v_N$  [Fig. 3(a)] that can be used in data-driven GMRES to simulate the permittivity distribution under consideration. To train the convolutional neural network, we consider two different loss functions:

- (a) *Projection loss function:* For the  $k^{\text{th}}$  training example, the projection loss  $l_{\text{proj}}^{(k)}$  is defined by the square of length of the simulated field  $f^{(k)}$  that is perpendicular to the space spanned by the vectors  $v_1, v_2 \dots v_N$  relative to the square of the length of  $f^{(k)}$ :

$$l_{\text{proj}}^{(k)} = \min_{f \in \mathcal{V}} \frac{\|f - f^{(k)}\|^2}{\|f^{(k)}\|^2} \quad (6)$$

Note that  $0 \leq l_{\text{proj}}^{(k)} \leq 1$ , with  $l_{\text{proj}}^{(k)} = 0$  indicating that  $f^{(k)}$  lies in the subspace  $\mathcal{V}$  and  $l_{\text{proj}}^{(k)} = 1$  indicating that  $f^{(k)}$  is orthogonal to  $\mathcal{V}$ .

- (b) *Residual loss function*: The residual loss function can remedy this issue. For the  $k^{\text{th}}$  training example, the residual loss  $l_{\text{res}}^{(k)}$  is defined by:

$$l_{\text{res}}^{(k)} = \min_{f \in \mathcal{V}} \frac{\|A^{(k)}f - b\|^2}{\|b\|^2} \quad (7)$$

where  $A^{(k)}$  is the sparse matrix corresponding to the operator  $\nabla \times \nabla \times -\omega^2 \varepsilon^{(k)}/c^2$  for the  $k^{\text{th}}$  training example. Similar to the projection loss,  $0 \leq l_{\text{res}}^{(k)} \leq 1$ , with  $l_{\text{res}}^{(k)} = 0 \Rightarrow f^{(k)} = [A^{(k)}]^{-1}b$  lies within the subspace  $\mathcal{V}$  and  $l_{\text{res}}^{(k)} = 1 \Rightarrow f^{(k)} = [A^{(k)}]^{-1}b$  is orthogonal to the subspace  $\mathcal{V}$  with respect to the positive definite matrix  $W^{(k)} = [A^{(k)}]^\dagger A^{(k)}$  (i.e.  $\langle v, f^{(k)} \rangle_{W^{(k)}} = v^\dagger W^{(k)} f^{(k)} = 0 \forall v \in \mathcal{V}$ ). However, unlike the projection loss function, the residual loss function is an unsupervised loss function i.e. the simulated electric fields for the structure are not required for its computation.

Since the residual loss function trains the neural network to minimize the residual directly, it accelerates data-driven GMRES significantly in the first few iterations as compared to the projection loss function. This can be clearly seen in Fig. 3(b), which shows histograms of the residual after 1 iteration and 100 iterations over the evaluation dataset. However, as is seen from the residual vs iteration number plots in Fig. 3(c), data-driven GMRES with the convolutional neural network trained using the residual loss function slows down after the first few iterations. It is not clear why this slow down happens, and is part of ongoing research.

We also note that the use of fields obtained with an effective index simulation as an input to the convolutional neural network significantly improves the performance of the convolutional neural network. Empirically, we observed that if the convolutional neural network was trained with the projection loss function but without the effective index fields as an input, the training loss would saturate at  $\sim 0.2$ , indicating that on an average approximately 20% of the simulated fields  $f^{(k)}$  lie outside the subspace  $\mathcal{V}$ . With a convolutional neural network that takes the effective index fields as an input in addition to the permittivity distribution, the training loss would saturate at a significantly lower value of  $\sim 0.05$  indicating that 95% of the simulated field  $f^{(k)}$  lies in the subspace  $\mathcal{V}$ .

Finally, we point out that the convolutional neural network needs to be trained only once before using it to accelerate GMRES over all the simulations in the evaluation dataset. The results presented in this paper were obtained using a convolutional neural network that required  $\sim 10$  hours to train, with the training distributed over  $\sim 8$  GPUs. However, once the neural network was trained, computing the vectors  $v_1, v_2 \dots v_N$  for a given structure (i.e. performing the feedforward computation for the neural network shown in Fig. 3(a)) is extremely fast, and the time take for this computation is negligible when compared to the iterative solve. For full three-dimensional simulations, we expect the iterative solve to be even slower than feedforward computations of most convolutional neural network and consequently believe that data-driven approaches can significantly accelerate electromagnetic simulations in practical settings.

**Benchmarks against data-free preconditioning techniques.** Accelerating the solution of linear system of equations is a problem that has been studied by the scientific computing community via data-free approaches for several decades. For iterative solvers in particular, using various preconditioners<sup>16</sup> to improve the spectral properties of the system of linear equations has emerged as a common strategy to solve this problem. In order to gauge the performance of the data-driven approaches outlined in this paper, we benchmark their performance against three stationary preconditioning techniques (Jacobi preconditioner, Gauss-Seidel preconditioner, Symmetric over-relaxation (SOR) preconditioner), incomplete LU preconditioner and a preconditioner designed specifically for finite-difference frequency domain simulations of Maxwell's equations<sup>17</sup> (see the supplement for residual vs. iteration plots of GMRES with different preconditioners on the evaluation dataset). Table 1 shows the results of these benchmarks:

1. *Training time*: While using data-driven GMRES augmented with principal components, the training time refers to the amount of time taken to compute the PCA vectors from the training data. While using data-driven GMRES augmented with the output of the CNN, training time refers to the amount of time taken to train the convolutional neural network by minimizing the projection loss function or the residual loss function over the training data. It can be noted that training is done only once, and the same PCA vectors or the same CNN are used for augmenting data-driven GMRES over all the structures in the evaluation dataset.
2. *Setup time*: The setup time refers to the amount of time taken to compute the augmenting vectors or the preconditioner for a given evaluation data sample. Note that there is no setup time while using principal components since they are computed once from the training data. The setup time for the convolutional neural network refers to the amount of time taken to perform the feedforward computation. The setup time for the preconditioners refers to the time taken for computing the preconditioners.
3. *GMRES time*: This refers to the amount of time taken to run GMRES iterations on the evaluation data sample to reduce the residual below a threshold  $r_{\text{th}}$ . The results in Table 1 are for  $r_{\text{th}} = 0.04$  which is chosen to be equal to the mean residual achieved by unpreconditioned GMRES after 100 iterations.
4. *Total solve time*: This is the amount of time taken for performing one simulation. It is calculated as a sum of the setup time (i.e. time required for calculating the augmenting vectors or the preconditioner) and the GMRES time. We do not include the training time in the solve time since the training is done exactly once for the entire evaluation dataset.
5. *Number of iterations*: This refers to the number of iterations for which GMRES needs to be performed to reduce the residual to below a threshold residual  $r_{\text{th}} = 0.04$ .



Method	Training time	Setup time	GMRES time	Total solve time	Number of iterations
GMRES	×	×	23.34 s	23.34 s	115.7
PCA: $N = 5$	7.63 s	×	0.87 s	0.87 s	6.1
PCA: $N = 10$	7.63 s	×	0.59 s	0.59 s	3.5
PCA: $N = 25$	7.63 s	×	0.62 s	0.62 s	2.1
PCA: $N = 50$	7.63 s	×	0.98 s	0.98 s	2.0
CNN: Projection loss	~10hr	0.18 s	13.46 s	13.52 s	64.5
CNN: Residual loss	~10hr	0.40 s	13.07 s	13.47 s	63.3
Jacobi	×	0.57 ms	39.30 s	39.30 s	153.8
Gauss seidel	×	17.59 ms	136.82 s	136.82 s	107.3
Ref. <sup>17</sup>	×	5.80 ms	24.87 s	24.87 s	118.7
SOR: $\omega = 0.25$	×	16.40 ms	170.17 s	170.17 s	132.0
SOR: $\omega = 0.50$	×	17.55 ms	152.21 s	152.21 s	120.2
SOR: $\omega = 0.75$	×	17.10 ms	140.56 s	140.56 s	111.7
SOR: $\omega = 1.25$	×	18.33 ms	138.67 s	138.67 s	107.5
SOR: $\omega = 1.50$	×	20.55 ms	169.73 s	169.73 s	120.4
SOR: $\omega = 1.75$	×	21.37 ms	349.27 s	349.27 s	235.9
ILU: Drop tol. = $10^{-1}$	×	6.62 s	13.90 s	20.52 s	66.2
ILU: Drop tol. = $10^{-2}$	×	7.01 s	4.59 s	11.60 s	28.6
ILU: Drop tol. = $10^{-3}$	×	9.29 s	0.74 s	10.03 s	3.0

**Table 1.** Benchmarking data-driven GMRES against data-free preconditioning techniques. The *training time* refers to the amount of time taken to compute the principal components or train the convolutional neural network using the training data. This is done only once, and the same principal components or trained neural network are used for augmenting GMRES over the evaluation data. The *setup time* is the time taken for computing the augmenting vectors or for computing the preconditioner for an evaluation data sample. The GMRES time is the time taken for running GMRES to reach a threshold residual of  $r_{th} = 0.04$  over the evaluation data. The *total solve time* is computed by summing the setup time and the GMRES time. The *number of iterations* shown are the number of GMRES iterations required to reduce the residual to below  $r_{th} = 0.04$ . The setup time, GMRES time and number of iterations shown in the table are averaged values computed over 50 randomly chosen evaluation data samples. The evaluation times reported are for a machine with 16 CPU cores and 60 GB RAM.

The numbers presented for setup time, GMRES time and number of iterations are average values obtained by simulating 50 randomly chosen evaluation data samples. We see from Table 1 that data-driven GMRES outperforms every preconditioning technique other than the incomplete LU preconditioner by at least an order of magnitude in terms of total solve time as well as the number of GMRES iterations performed. The best case performance of these preconditioners is only slightly better than unpreconditioned GMRES. Preconditioning partial differential equations describing wave propagation is known to be a difficult problem due to the solution of the wave equations being delocalized over the entire simulation domain even when the source has compact spatial support<sup>15</sup>. We also note that while the incomplete LU preconditioner does provide a notable speedup over unpreconditioned GMRES, the PCA-based data-driven GMRES outperforms it in terms of the total solve time and CNN-based data-driven GMRES performs comparably to it. Clearly, these benchmarks indicate that the data-driven approaches introduced in this paper have the potential to provide a scalable alternative to data-free preconditioning techniques for accelerating electromagnetic simulations.

## Conclusion

In conclusion, we present a framework for accelerating finite difference frequency domain (FDFD) simulations of Maxwell's equations using data-driven models that can exploit simulations of correlated permittivity distributions. We analyze two data-driven models, based on principal component analysis and a convolutional neural network, to accelerate these simulations, and show that these models can unlock an orders of magnitude acceleration over data-free solver. Such data-driven methods would likely be important in scenarios where a large number of simulations with similar permittivity distributions are performed e.g. during a gradient-based optimization of a photonic device.

## Methods

**Dataset.** The grating splitters are designed to reflect an incident waveguide mode at  $\lambda = 1.4 \mu\text{m}$  and transmit an incident waveguide mode at  $\lambda = 1.55 \mu\text{m}$  using a gradient-based design technique similar to that used for grating couplers<sup>3</sup>. Different devices in the dataset are generated by seeding the optimization with a different initial structure. Note that all the devices generated at different stages of the optimization are part of the dataset. Consequently, the dataset not only has grating splitters that have a discrete permittivity distribution (i.e. have only two materials – silicon and silicon oxide), but also grating splitters that have a continuous distribution (i.e. permittivity of the grating splitter can assume any value between that of silicon oxide and silicon). Moreover, the dataset has poorly performing grating splitters (i.e. grating splitters generated at the initial steps of the optimization procedure) as well as well-performing grating splitters (i.e. grating splitters generated towards the end of the optimization procedure). Our dataset has a total of ~30,000 examples, which we split into a training data set (75%) and an evaluation dataset (25%).

**Implementation of data-driven models.** Both the data-driven models (PCA and CNN) were implemented using the python library Tensorflow<sup>18</sup>. Any complex inputs (e.g. effective index fields) to the CNN were fed as an image of depth 2 comprising of the real and imaginary parts of the complex input. We use the ADAM optimizer<sup>19</sup> with a batch size of 30 for training the convolutional neural network — it required ~10,000 steps to train the network.

Received: 22 February 2019; Accepted: 2 December 2019;

Published online: 23 December 2019

## References

1. Piggott, A. Y. *et al.* Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer. *Nature Photonics* **9**, 374–377 (2015).
2. Piggott, A. Y. *et al.* Inverse design and implementation of a wavelength demultiplexing grating coupler. *Scientific reports* **4**, 7210 (2014).
3. Su, L. *et al.* Fully-automated optimization of grating couplers. *Optics express* **26**, 4023–4034 (2018).
4. Parks, M. L., De Sturler, E., Mackey, G., Johnson, D. D. & Maiti, S. Recycling krylov subspaces for sequences of linear systems. *SIAM Journal on Scientific Computing* **28**, 1651–1674 (2006).
5. Wang, S., Sturler, E. D. & Paulino, G. H. Large-scale topology optimization using preconditioned krylov subspace methods with recycling. *International journal for numerical methods in engineering* **69**, 2441–2468 (2007).
6. Bar-Sinai, Y., Hoyer, S., Hickey, J. & Brenner, M. P. Data-driven discretization: a method for systematic coarse graining of partial differential equations. *arXiv preprint arXiv:1808.04930* (2018).
7. Han, J., Jentzen, A. & Weinan, E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences* 201718942 (2018).
8. Golub, G. H. & Van Loan, C. F. *Matrix computations*, vol. 3 (JHU Press, 2012).
9. Peurifoy, J. *et al.* Nanophotonic particle simulation and inverse design using artificial neural networks. *Science advances* **4**, eaar4206 (2018).
10. Liu, D., Tan, Y., Khoram, E. & Yu, Z. Training deep neural networks for the inverse design of nanophotonic structures. *ACS Photonics* **5**, 1365–1369 (2018).
11. Shin, W. *3D finite-difference frequency-domain method for plasmonics and nanophotonics*. Ph.D. thesis, Stanford University (2013).
12. Hicken, J. E. & Zingg, D. W. A simplified and flexible variant of gcrot for solving nonsymmetric linear systems. *SIAM Journal on Scientific Computing* **32**, 1672–1694 (2010).
13. Hammer, M. & Ivanova, O. V. Effective index approximations of photonic crystal slabs: a 2-to-1-d assessment. *Optical and quantum electronics* **41**, 267–283 (2009).
14. Jolliffe, I. Principal component analysis. In *International encyclopedia of statistical science*, 1094–1096 (Springer, 2011).
15. Ernst, O. G. & Gander, M. J. Why it is difficult to solve helmholtz problems with classical iterative methods. In *Numerical analysis of multiscale problems*, 325–363 (Springer, 2012).
16. Saad, Y. *Iterative methods for sparse linear systems*, vol. 82 (SIAM, 2003).
17. Shin, W. & Fan, S. Simulation of phenomena characterized by partial differential equations (US patent, 2014).
18. Abadi, M. Tensorflow: learning functions at scale. In *Acm Sigplan Notices*, vol. 51, 1–1 (ACM, 2016).
19. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

## Acknowledgements

R.T. acknowledges funding from Kailath graduate fellowship. This work has been supported Google Inc, grant number 134920.

## Author contributions

R.T. and L.S. conceived the idea, R.T. and J.L. designed and implemented the numerical experiments, M.F.S. and J.V. supervised the project. All authors contributed to the discussion and analysis of results.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** is available for this paper at <https://doi.org/10.1038/s41598-019-56212-5>.

**Correspondence** and requests for materials should be addressed to R.T.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019