Physics-guided Training of Neural Electromagnetic Wave Simulators with Time-reversal Consistency

Charles Dove, Jatearoon Boondicharern, Laura Waller
Department of Electrical Engineering and Computer Science
University of California, Berkeley
Berkeley, CA
{charles_dove,jboondicharern,waller}@berkeley.edu

Abstract

Conventional electromagnetic wave simulators often have long simulation times, so are not suitable for computational imaging and photonic inverse problems (e.g. end-to-end design, iterative reconstruction) that require evaluating the forward model many times. Electromagnetic wave simulators based on neural networks promise speed improvements of several orders-of-magnitude, but standard supervised training approaches have difficulty fitting the true physics. Physics-informed approaches help, but existing residual-based methods use only local information and must be used in conjunction with standard supervised loss. In this work, we introduce Time Reversal Consistency (TReC), a new physics-based training method based on the time reversibility of Maxwell's equations. TReC uses a time-reversed, differentiable finite-difference simulator to compare neural network predictions with a known initial condition. TReC provides both global physics guidance and supervision in a single function. When trained only on randomized scatterers, we find that networks trained with TReC generalize well to a range of arbitrary structured media. We validate the method on the inverse design of a set of angle-to-angle couplers, addressing almost two magnitudes more parameters than previous methods, and find that the design quality corresponds closely with designs based on a conventional simulator while requiring 5% of the design time.

1 Introduction and Background

Electromagnetic (EM) wave simulators based on Maxwell's equations provide a detailed quantitative understanding of light-matter interactions. Understanding and shaping these interactions is fundamental to various fields, from computational microscopy [1] to photonics [2] and radar [3]. Traditional simulation methods like the finite difference time domain (FDTD) [4] method are accurate but computationally demanding, with their compute and runtime requirements growing quickly with the system's spatial size. This is especially limiting for EM wave inverse problems, which require hundreds or thousands of consecutive simulations to solve an iterative optimization, along with an intensive gradient calculation process for each. There is therefore a pressing need for faster, more efficiently differentiable EM wave simulators.

Neural network (NN)-based simulators, also called neural simulators, have the potential to provide fast EM field predictions for a given source and refractive index (RI) configuration, along with gradient information by backpropagation. Building a scalable, accurate neural EM wave simulator, however, has proven challenging. The problem is most commonly approached with supervised direct prediction (SDP) [5, 6, 7, 8, 9, 10], in which the network is trained via standard supervised learning to directly predict a final, converged EM field from the RI configuration. While fast, this method requires an amount of ground-truth training data which scales in most cases exponentially [11] with the number of input parameters. Here, the number of input parameters refers to the number of pixels with different RI values that are modeled within the spatial area simulated. For instance, a 20 input-parameter SDP network has been found to require a pregenerated dataset of 750,000

37th Conference on Neural Information Processing Systems (NeurIPS 2023).

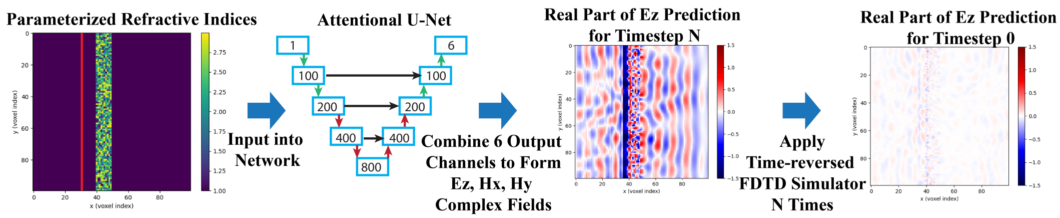

Figure 1: Time Reversal Consistency (TReC) training method. Parameterized refractive index (RI), with source location in red, are input into an attentional U-Net. Green arrows indicate 3x3 convolutions and red indicate 3x3 convolutions with a subsequent attentional layer, with 2x2 downsampling/upsampling indicated by arrow direction and black skip connections. Channel dimensions are noted in the blue boxes. The network outputs predictions of E_z, H_x, and H_y complex-fields ($Re(E_z)$ shown) for timestep N. These are differentiably propagated backward in time to timestep 0 and compared to the known initial condition to form the loss function, L_{TReC}.

examples [6] to be suitable for inverse design applications, while a 1,000 input-parameter network has been projected to require an intractable one billion examples [11]. Additionally, gradients produced by backpropagation through SDP methods tend to be low quality, often necessitating the use of a problem-specific auxiliary adjoint neural network [5] which requires its own distinct training dataset.

Physics-informed loss functions [5, 12, 7] have been used to improve network fitting by minimizing a wave equation-based residual. While these provide a meaningful fitting benefit compared to pure supervised learning, they are based on spatially local residuals, which do not capture nonlocal physically-relevant correspondences, contributing to a loss landscape which is challenging to optimize. When used alone, NNs trained with these residuals have difficulty converging to accurate solutions. As such, physics-informed loss functions are generally only used as a form of regularizer added to a supervised loss.

In this work, we introduce Time Reversal Consistency (TReC), a new physics-based training method for direct prediction networks based on the time reversibility of Maxwell's equations. Unlike physics-informed loss functions based on local residual minimization, TReC combines both global physics guidance and direct supervision, comparing network outputs through a differentiable simulation to the known initial conditions of a simulation. We find that using TReC alone, trained on systems with randomized input parameters, allows efficient, highly accurate training of NNs for complex, highly-scattering systems with thousands of input parameters. We additionally find that, when trained purely on random media, the TReC network generalizes well to structured media. Instead of using an auxiliary adjoint network to calculate gradients, we find that gradients calculated by backpropagation through a TReC NN trained purely on randomized media are suitable for inverse problems. To demonstrate this, we show inverse design of 1000-parameter angle-to-angle coupling devices, which we believe to be nearly two orders-of-magnitude more parameters than the state of the art of 20 [6]. Instead of requiring 1 billion independent training simulations to create a NN capable of this, TReC produces results close to those provided by FDTD with only 700,000 training simulations.

2 Results

2.1 Time reversal consistency

We developed TReC, diagrammed in Fig. 1, which works by taking the prediction of a direct prediction neural network tasked with predicting the final state of an N-timestep simulation, then applying a differentiable time-reversed simulator N times to this prediction, producing a corresponding state at timestep 0. Because the EM state at time 0 is known–and in most situations zero everywhere–this reversed prediction can then be directly compared to the initial state in a loss function, then backpropagated through the entire physical interaction and NN, capturing global physical dependencies. For a time-reversed FDTD step f^{-1}, EM field prediction at time N $\hat{\mathbf{y}}_N$, known EM initial condition $\mathbf{y}_0$, and Mean Squared Error (MSE) loss function, the TReC loss function can be stated as

$$L_{TReC}(\mathbf{y}_0, f^{-N}(\hat{\mathbf{y}}_N)) = \mathbf{MSE}(E_0, \hat{E}_0) + \mathbf{MSE}(H_0, \hat{H}_0) \tag{1}$$

where each $\mathbf{y}$ contains corresponding E and H fields. $\mathbf{y}_0 = E_0 = H_0 = \mathbf{0}$ in most practical cases.

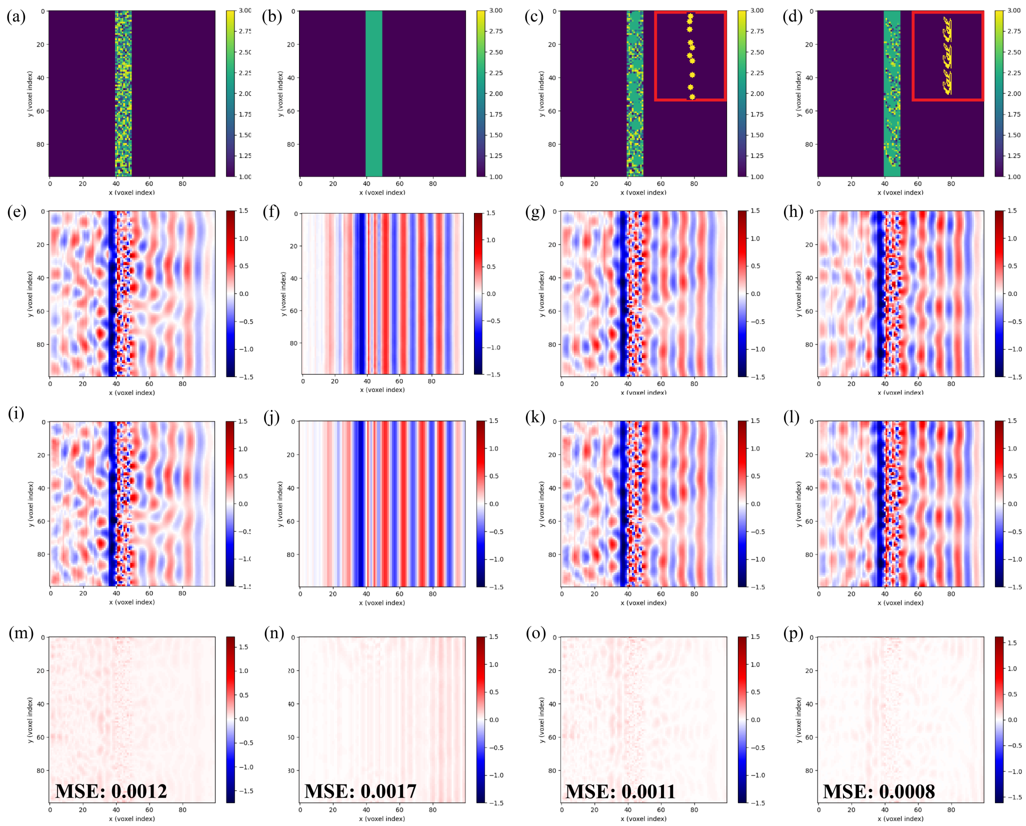

Figure 2: (a-d) RIs for four different objects: random, monolithic, microsphere, and Cal logo, with masks inset indicating microsphere and logo locations. (e-h) Real part of NN prediction of E_z for each object (i-l) Real part of FDTD ground truth of E_z for each object (m-p) Absolute difference map between prediction and ground truth, with MSE inset.

2.2 Architecture and training

While any network architecture can be used with TReC, we here use a 4-level convolutional U-net, with 3x3 kernels for each layer, two-layer residual blocks at each down and up-sampling level, for a total of 16 layers, and 4-head attentional layers [13] on the innermost two blocks for both the downsampling and upsampling side. The NN takes a grid of RIs as input and outputs 6 total channels, which form the real and imaginary parts of the E_z, H_x, and H_y fields.

Training is accomplished using a purpose-built differentiable time-reversed FDTD simulator built in the Python library Pytorch [10]. For demonstration, we consider a 100x100 region of rectangular voxels, each with side length of 45 nm, with an included 100x10 volumetric scatterer left-illuminated by a plane wave of 450 nm blue light, with periodic boundary conditions connecting the top and bottom edges. We choose the number of timesteps $N = 120$. Each of the 1000 scatterer voxels is taken to be an independent parameter indicating the refractive index at that point with range $[1, 3]$. For training, a batch of random scatterers are just-in-time generated by drawing parameters from $\mathcal{U}(1, 3)$. This batch of random scatterers is then fed through the NN, and L_{TReC} is applied to the output using the differentiable FDTD simulator to compute f^{-N}. Network weights are updated through backpropagation, and the process is repeated until satisfactory convergence. We use the Adam optimizer with learning rate 1×10^{-4} and a batch size of 20. Validation and accuracy quantification is performed by comparing network outputs at time N to a corresponding N-step forward FDTD simulation, using MSE. All training is performed on a single Nvidia A6000 GPU.

2.3 Relationship between TReC loss and supervised loss

While providing physics guidance, TReC also has a close relationship with supervised loss functions. This is due to the linearity of Maxwell's equations and energy conservation. Intuitively, the NN prediction $\hat{\mathbf{y}}_N$ can be decomposed into the exact field $\mathbf{y}_N$ and an error field $\mathbf{e}_N$. Applying the time-reversed simulator N times exactly cancels $\mathbf{y}_N$ but, assuming nonabsorptive materials, leaves

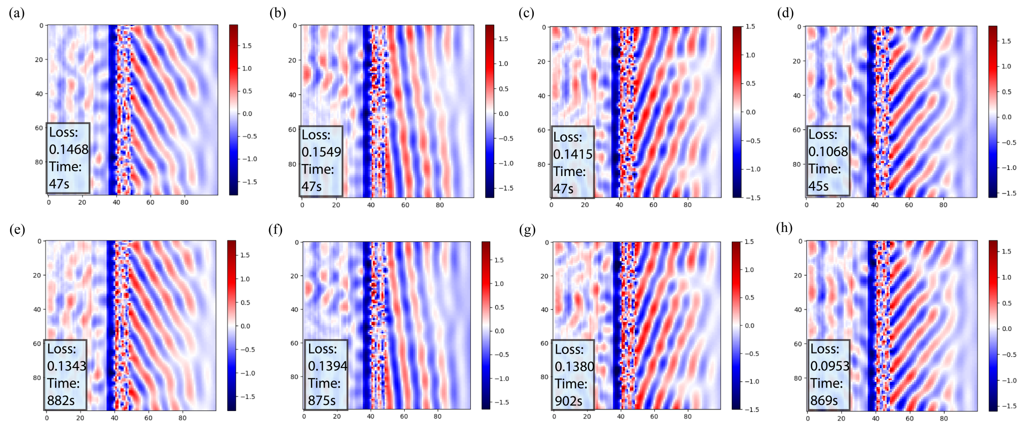

Figure 3: Angle-to-angle coupler designs. (a-d) $Re(E_z)$ field for the TReC designs with design angles of -45, -10, 30, and 60 degrees, respectively. The design loss, the MSE between the output wavefront profile and the ideal wavefront profile, is inset, as well as runtime. (e-h) $Re(E_z)$ field of FDTD-based designs for the same angles, with MSE design loss and runtime inset.

the same amount of energy in $\mathbf{e}_0$ as $\mathbf{e}_N$. Neglecting the influence of nonuniform permittivities, which have a linear reweighting effect which we find empirically has minor impact on fitting when incorporated into the loss, this $\mathbf{e}_0$ energy corresponds mathematically with L_{TReC}, and it can be shown that the energy of $\mathbf{e}_N$ corresponds with a simple MSE-based supervised loss L_{sup}. As such, $L_{TReC} \propto L_{sup}$. A derivation of this property is provided in Appendix A.1.

2.4 Generalization to structured media

While we train the TReC NN solely on randomized systems, we find that the NN generalizes well to a range of arbitrary structured scenerios. To demonstrate this, we first train the NN to full convergence, taking about 8 hours on a single Nvidia A6000 GPU and reaching a validation MSE of 0.00119 after seeing 700,000 total examples. We then test the NN on a range of arbitrary parameter configurations, shown in Fig. 2. We find strong performance for random scatterers (the validation dataset), monolithic media, microspheres embedded in scattering media, and scattering logos in monolithic media.

2.5 Inverse design of 1000-parameter angle-to-angle couplers

A primary application of TReC is in inverse problems, such as those found across computational imaging and photonic inverse design. To demonstrate TReC's efficacy and correspondence with FDTD in such applications, we consider the design of four angle-to-angle couplers, with each coupling a given plane wave to a chosen, arbitrary angle. We accomplish the designs by comparing a device's E_z field profile directly above the right surface with the desired output field profile via MSE, then updating the device RIs via backpropagation and an Adam optimizer with a learning rate of 1×10^{-4}. We start from the randomized 1000-parameter scatterer used in previous sections, left-illuminated by a normally-incident 450 nm blue plane wave, and accomplish the designs by continuously tuning the RI for each parameter. RIs are constrained via a proximal step to lie in the range $[1, 3]$. For the designs, we use, alternately, the noise-trained NN described in the previous section and the ground-truth FDTD simulator, backpropagating through each to calculate gradients for the input parameters. With results shown in Fig. 3, we use FDTD to compare the designs produced after 1000 iterations for angle-to-angle gratings with design angles of -45, -10, 30, and 60 degrees, respectively. RIs for each design are included in Appendix A.2. We find that the TReC designs perform with a design loss within 11% of the FDTD-based designs, while requiring only 5% of the total design time.

3 Conclusion

We introduced TReC, a new method for physics-guided training of neural EM wave simulators. TReC provides both global physics guidance and supervision, strong generalization capacity when trained on randomized systems, and high quality parameter gradients suitable for many-parameter inverse design tasks. TReC provides a practical path to many-parameter neural EM wave simulators for inverse problems, with applications ranging from real-time reconstructive imaging to freeform photonic design.

References

[1] Michael Chen, Laura Waller, David Ren, Shwetadwip Chowdhury, and Hsiou-Yuan Liu. Multi-layer born multiple-scattering model for 3d phase microscopy. *Optica, Vol. 7, Issue 5, pp. 394-403*, 7:394–403, 5 2020.

[2] Eli Yablonovitch and Andrew Michaels. Leveraging continuous material averaging for inverse electromagnetic design. *Optics Express, Vol. 26, Issue 24, pp. 31717-31737*, 26:31717–31737, 11 2018.

[3] Sixin Liu, Zhaofa Zeng, and Li Deng. Fdtd simulations for ground penetrating radar in urban applications. *Journal of Geophysics and Engineering*, 4:262–267, 9 2007.

[4] Kane S. Yee. Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media. *IEEE Transactions on Antennas and Propagation*, 14:302–307, 1966.

[5] Mingkun Chen, Robert Lupoiu, Chenkai Mao, Der Han Huang, Jiaqi Jiang, Philippe Lalanne, and Jonathan A. Fan. High speed simulation and freeform optimization of nanophotonic devices with physics-augmented deep learning. *ACS Photonics*, 9:3110–3123, 9 2022.

[6] Dianjing Liu, Yixuan Tan, Erfan Khoram, and Zongfu Yu. Training deep neural networks for the inverse design of nanophotonic structures. *ACS Photonics*, 5:1365–1369, 4 2018.

[7] Joowon Lim and Demetri Psaltis. Maxwellnet: Physics-driven deep neural network training based on maxwell's equations. *APL Photonics*, 7, 1 2022.

[8] Wei Ma, Zhaocheng Liu, Zhaxylyk A. Kudyshev, Alexandra Boltasseva, Wenshan Cai, and Yongmin Liu. Deep learning for the design of photonic structures. *Nature Photonics 2020 15:2*, 15:77–90, 10 2020.

[9] Sunae So, Trevon Badloe, Jaebum Noh, Junsuk Rho, and Jorge Bravo-Abad. Deep learning enabled inverse design in nanophotonics. *Nanophotonics*, 9:1041–1057, 5 2020.

[10] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury Google, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf Xamla, Edward Yang, Zach Devito, Martin Raison Nabla, Alykhan Tejani, Sasank Chilamkurthy, Qure Ai, Benoit Steiner, Lu Fang Facebook, Junjie Bai Facebook, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32, 2019.

[11] Jiaqi Jiang, Mingkun Chen, and Jonathan A. Fan. Deep neural networks for the evaluation and design of photonic devices. *Nature Reviews Materials 2020 6:8*, 6:679–700, 12 2020.

[12] Yingheng Tang, Jichao Fan, Xinwei Li, Jianzhu Ma, Minghao Qi, Cunxi Yu, and Weilu Gao. Physics-informed recurrent neural network for time dynamics in optical resonances. *Nature Computational Science 2022 2:3*, 2:169–178, 3 2022.

[13] Ashish Vaswani, Google Brain, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017.

A Appendix

A.1 Derivation of TReC loss, supervised loss correspondence

To demonstrate the correspondence of TReC, evaluated at timestep 0, to a supervised loss evaluated at timestep N, we consider decomposing the EM prediction at time N via linearity $\hat{\mathbf{y}}_N = \mathbf{y}_N + \mathbf{e}_N$ into exact EM field $\mathbf{y}_N$ and remaining error EM field $\mathbf{e}_N$. Neglecting nonuniform permittivities and constants, the error field has a total energy U_{eN}

$$\begin{gathered} U_{eN} \propto \int_A \left((\mathbf{E}_{\hat{y}N} - \mathbf{E}_{yN})^2 + (\mathbf{H}_{\hat{y}N} - \mathbf{H}_{yN})^2 \right) \, dA \propto \\ \mathbf{MSE}(\mathbf{E}_{yN}, \mathbf{E}_{\hat{y}N}) + \mathbf{MSE}(\mathbf{H}_{yN}, \mathbf{H}_{\hat{y}N}) = L_{sup} \end{gathered} \tag{2}$$

proportional to the MSE-based supervised loss L_{sup}. Similarly, supposing $\mathbf{y}_0 = \mathbf{0}$, the time-reversed error field $\mathbf{e}_0 = f^{-N}(\hat{\mathbf{y}}_N) = \mathbf{0} + f^{-N}(\mathbf{e}_N) = \hat{\mathbf{y}}_0$ has energy

$$U_{e0} \propto \int_A \left((\mathbf{E}_{\hat{y}0})^2 + (\mathbf{H}_{\hat{y}0})^2 \right) \, dA \propto \mathbf{MSE}(\mathbf{0}, \mathbf{E}_{\hat{y}0}) + \mathbf{MSE}(\mathbf{0}, \mathbf{H}_{\hat{y}0}) = L_{TReC} \tag{3}$$

Applying the time-reversed simulator N times will by construction exactly cancel the EM field of y_N. In the absence of absorptive materials or boundaries, however, the energy of eN will be retained over this process. As such $U_{e0} = U_{eN}, Uy0 = 0$ and therefore

$$L_{sup} \propto L_{TReC} \tag{4}$$

A.2 Angle-to-angle coupler refractive indices

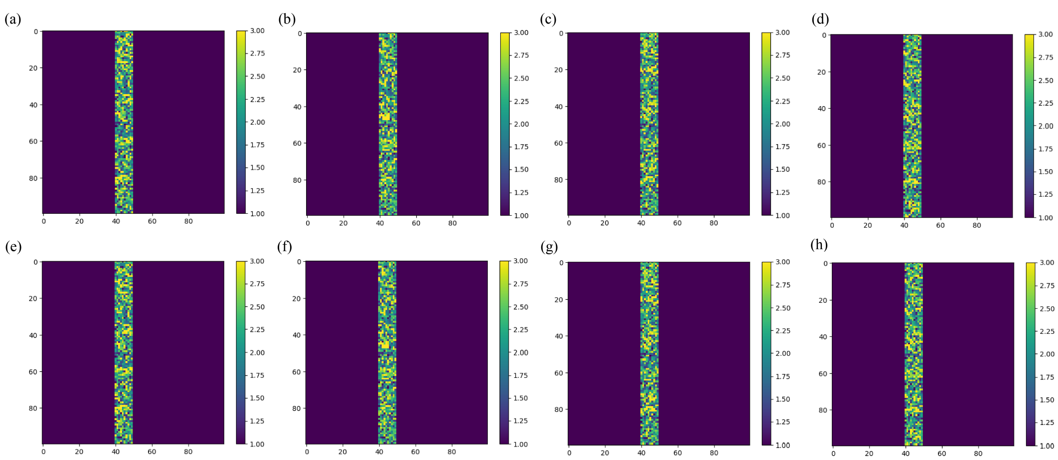

Figure 4: Refractive index configurations for each angle to angle coupler design, as described in Section 2.5. (a-d) Refractive indices for the TReC designs with design angles of -45, -10, 30, and 60 degrees, respectively. (e-h) Refractive indices of the FDTD-based designs for the same angles.