

Inverse Design of Unitary Transmission Matrices in Silicon Photonic Coupled Waveguide Arrays Using a Neural Adjoint Model

Thomas W. Radford,* Peter R. Wiecha, Alberto Politi, Ioannis Zeimpekis, and Otto L. Muskens*



Cite This: *ACS Photonics* 2025, 12, 1480–1493



Read Online

ACCESS |



Metrics & More




Article Recommendations



Supporting Information

ABSTRACT: The development of low-loss reconfigurable integrated optical devices enables further research into technologies including photonic signal processing, analogue quantum computing, and optical neural networks. Here, we introduce digital patterning of coupled waveguide arrays as a platform capable of implementing unitary matrix operations. Determining the required device geometry for a specific optical output is computationally challenging and requires a robust and versatile inverse design protocol. In this work we present an approach using high speed neural network surrogate-based gradient optimization, capable of predicting patterns of refractive index perturbations based on switching of the ultralow loss chalcogenide phase change material, antimony triselenide (Sb_2Se_3). Results for a 3×3 silicon waveguide array are presented, demonstrating control of both amplitude and phase for each transmission matrix element. Network performance is studied using neural network optimization tools such as data set augmentation and supplementation with random noise, resulting in an average fidelity of 0.94 for unitary matrix targets. Our results show that coupled waveguide arrays with perturbation patterns offer new routes for achieving programmable unitary operators, or Hamiltonians for quantum simulators, with a reduced footprint compared to conventional interferometer-mesh technology.

KEYWORDS: inverse design, deep learning, neural adjoint, programmable photonic devices, phase change materials, Sb_2Se_3



INTRODUCTION

Integrated photonics offers a platform for the miniaturization of optical devices and systems, yielding increased stability, greatly reduced size and complexity compared to traditional optical systems. Integrated photonic technologies have been successfully demonstrated across a wide range of domains including vector-matrix multiplication,^{1,2} quantum simulation,^{3,4} signal processing⁵ and the development of high speed neural networks.⁶ While application-specific photonic integrated circuits require a bespoke design-fabrication cycle for each variation in functionality, there is an emerging interest in platforms that can be reprogrammed after fabrication to provide fine-tuning, diversification, or entirely new methods of deployment.⁷ Reprogrammable devices may also have the ability to dynamically tailor their optical output,⁸ and use cases were realized across a number of fields including post fabrication device processing,⁹ optical switches¹⁰ and optical signal compensation.¹¹ One field in particular looking to utilize reconfigurable technologies is photonic computing^{6,12–14} where the development of next generation optical devices rely on the construction of structures with rapidly configurable output characteristics allowing for high speed information processing.

Commonly used reconfigurable technologies such as micro-heaters^{15,16} and micro electro-mechanical devices (MEMS)^{17–19} either introduce additional complex fabrication steps or rely on the coupling of external regulating electronics onto the photonic chip. Such approaches occupy valuable chip space and impose power constraints, while simultaneously introducing undesirable heat and noise into experimental systems. An alternative method for the production of reconfigurable devices involves using nonvolatile optical phase change materials (PCMs) such as antimony triselenide (Sb_2Se_3).²⁰ PCM based approaches can offer ultralow loss, compact and passive programming capabilities, while maintaining compatibility with a wide range of optical devices across telecommunications wavelengths.^{21–25} PCMs such as Sb_2Se_3 are able to exist in multiple nonvolatile solid states which can be programmed using optical or electrical pulses. The dielectric functions of amorphous and crystalline states of

Received: October 23, 2024

Revised: January 29, 2025

Accepted: January 30, 2025

Published: February 12, 2025






Figure 1. (a) Artistic view of the reconfigurable coupled waveguide device. An array of coupled silicon rib waveguides (here three) is covered by a thin PCM layer (here, Sb_2Se_3), which can be locally switched between crystalline and amorphous state using a laser writing setup. This is parametrized in a pixel pattern, in this figure, 500 nm large pixels along $50\mu\text{m}$ of the waveguides. The input and output ports of the waveguide array can be individually accessed by optical gratings. (b) varFDTD simulations of the electric field intensity maps along the coupled waveguide array for three different input ports. The pattern is designed to correspond to an antidiagonal exchange matrix with equal phases for each output. (c) Full schematic of the inverse design pipeline used to find the optimal pattern shown in (b).

PCMs differ significantly from each other, which depending on the nature of the chemical bonds can result in large order changes in refractive index.²⁶ An approach using direct optical writing allows the addition of reconfigurability to previously fabricated devices using minimal overheads in fabrication and regulation.²⁷ Endurance in programmable devices using PCM technologies is seeing improvements²⁸ with repeated refractive index modulation of $>10^6$ switching cycles recently being demonstrated in thin-films of Sb_2Se_3 .^{29,30}

The relationship between input and output modes of any linear optical element can be described by a complex transmission matrix. Through careful design of multiport devices, we are therefore able to produce an analogue for any given $n \times m$ matrix so long as a fabrication process can accommodate its production. Integrated approaches using continuously coupled devices offer an interesting platform whereby a device can be specifically designed to implement any arbitrary transmission matrix, confined within the bounds of a micron-scale structure.³¹ Interference based devices with the addition of individual scattering sites are a common platform for the implementation of arbitrary optical transmission matrices,^{32–37} as well as devices consisting of reconfigurable interferometer meshes.^{8,12,13,38}

In this work, we explore an alternative approach, based on coupled waveguide arrays (CWG) covered with a thin film of PCM programmed with a nanoscale perturbation pattern.^{39–45} With focused laser writing, as illustrated in Figure 1a, it is possible to reversibly modulate the local refractive index of individual pixels on the device surface with diffraction limited spatial resolution.^{23,27} This pattern facilitates modulation of the coupling coefficients between neighboring waveguides, allowing full control over the devices transmission properties. As light propagates through a part of the device containing a switched pixel, it develops a phase shift relative to light traveling in an unswitched region due to the local effective

index contrast. Using a large enough number of these small effects at each pixel, enables full control over both the waveguide coupling behavior, as well as the phase of the light at output ports across a full 2π range.

In contrast to microheater or interferometer mesh-based devices, the perturbation pixel patterns in our CWG platform are correlated to the device's transmission matrix in a highly complex manner. A crucial step to develop this technology must therefore be the development of a powerful inverse design protocol, used to find a pixel pattern that implements a specific, complex transmission matrix. Significant research has been undertaken to develop a variety of inverse design processes for nanophotonics.^{46,47} Modern approaches make use of techniques such as Bayesian optimization, Genetic algorithms and the utilization of automatic differentiation, with among the most popular methods those which carry out inverse design through topology optimization⁴⁸ and the integration of deep-neural networks. Topology optimizations are very powerful⁴⁹ and have found applications across a wide range of disciplines, from aerospace engineering^{50,51} and medical research,^{52,53} to photonics.^{54–58} Such optimizations typically result in complex-shaped structures with fine features which can prove difficult to fabricate, in some cases becoming unrealistic to produce at all.⁵⁹ Gradient based topology approaches often face practical limitations as gradient optimizations require smooth variables, but their design problem is typically categorized into a binary basis, where material either is present or absent.

It should be noted that while there do exist nongradient based optimization methods, many of these have been demonstrated as intractable for real world practical systems with large degrees of freedom.^{46,60} The most popular current alternative is the integration of artificial intelligence and machine learning into the device design pipeline, with

increasing attention being paid to the creation of an AI-topological hybrid approach.^{61–63}

In this work, we demonstrate an inverse design tool chain using a gradient based optimization with a deep learning surrogate forward model,⁶⁴ as well as a geometry reparametrization using a Wasserstein generative adversarial network (WGAN). Using this high speed technique we build on previous works in which similar methods are employed for the design and development of individual device elements, such as the height, orientation and shape of elliptic meta atom pillars,^{65,66} as well as in more complex free-space design tasks.⁶⁷

■ NEURAL ADJOINT INVERSE DESIGN APPROACH

Device Model. To demonstrate the capabilities of our inverse design method, we consider the case of an array of three coupled silicon rib waveguides, implementing an arbitrary 3×3 transmission matrix. In our simulations, these waveguides are covered by crystalline Sb_2Se_3 , which may then be switched locally, for example experimentally by using direct laser writing.^{23,27,68} Experimentally it is easier to produce a narrow amorphization laser profile, meaning greater spatial precision can be achieved by using pixel maps with a crystalline background. We consider thin films of PCM, which have been shown experimentally as a feasible geometry,²⁷ allowing a consistent and repeatable index modulation to be applied.

Our coupled waveguide device model is based on a standard silicon on insulator rib waveguide design, a schematic of the device is shown in Figure 1a, with detailed features as described in the Methods section. A 30 nm thick Sb_2Se_3 PCM layer on top of the CWG provides a switchable index contrast between crystalline and amorphous states of $\Delta n_{\text{Sb}_2\text{Se}_3} = 0.77$. Pixels are defined on a 500 nm square grid over the coupling region's area, consisting of a total of 96 pixels for each of the three waveguides (288 pixels in total).

Coupled waveguide devices are simulated using Variational FDTD (varFDTD) as described in the Methods section. This engine collapses a 3D device structure into an effective 2D simulation by probing the effective index experienced in slices across the simulation mesh. Figure 1b shows the electric field intensity of a typical patterned device that was designed for a transmission matrix equal to the inverse identity matrix. After simulation, the device transmission matrix is extracted by evaluating solutions for light coupled into each of the three input waveguides, where for all input modes the complex field at the peak amplitude of each output is recorded and constructed into a complex matrix.

Inverse Design Goals and Training Data Generation.

The choice of training data is a particular important aspect of the model under study. Similar to previous studies,³⁶ we use preconditioned data by specifying a target output profile for a single input port. In our studies we found that it is not constructive to use a training set of purely random patterns,³⁶ because random pixel patterns typically do not co-operate in providing an effective state conversion between input and output, as discussed in more detail later. For an acceptable performance of the deep learning based design method, a more directed approach to training data generation is required, one that from the start aims toward the later inverse design task.

To this end, we first implement a brute force forward iterative optimization algorithm in the training data generation as described in more detail in the Methods section and

Supporting Information, Figure S5. Iterative optimization is used to optimize the output for a single random input to a specific splitting ratio by switching individual pixels. This process effectively optimizes one single row of the matrix. At each evaluation step where a pixel is found to improve the objective, the transmission of the other input ports is also evaluated and the entire pattern and its associated complex transmission matrix are saved to the data set. In this way the data set is populated with a large variety of pixel numbers and sufficient variance in the device performance to build a balanced data set. A detailed description of this process as well as a specific example can be found in the Supporting Information, Figure S5. Training data generation was carried out to create an initial data set containing 42812 pixel patterns alongside their associated complex transmission matrices, where augmentation exploiting the device symmetry upon vertical reflection is used to double the size of the data set. Once a database of preoptimized pixel patterns and their associated transmission matrices is created, we repeat this process to create a separate validation data set used to verify that the network is sufficiently generalized and may accurately predict transmission matrices for patterns from outside of the data set it was trained on. This validation set contains a further 6564 patterns after augmentation.

Gradient Based Inverse Design Using a Deep Learning Surrogate. The forward predictor used here is a ResNeXt encoder-decoder convolutional neural network,⁶⁹ a detailed schematic of this forward network model can be found in the Supporting Information, Figure S1. While such purely data-based surrogate models are known to work accurately in the interpolation regime, it is also well-known that extrapolation to geometries that differ from those in the data set is usually very error-prone.^{70,71} Thus, a mechanism is necessary to constrain the allowed geometries during optimization to pixel patterns which the forward model is capable of faithfully predicting. Typically this is achieved through the introduction of a boundary loss term in the design optimization.⁶⁵ However, in free-form optimizations such a constraint is difficult to formulate.

To avoid the optimization converging to the extrapolation regime of the forward network, we use here a further deep-learning based approach. Integrating a Wasserstein generative adversarial network with gradient penalty⁷² (WGAN-GP), we develop a learned reparametrization of the pixel patterns geometry from the training data set. Details of the WGAN-GP architecture are shown in the Supporting Information, Figure S2. The key idea of this architecture is, that during training, the WGAN-GP develops a mapping of the pixel pattern geometries into a compact and continuous latent space, in which the pattern geometries of the training data are normally distributed around the mean value $\mu_z = 0$ and with a variance of $\sigma_z = 1$. The latent vector representations of the geometries in the data set being normally distributed in the latent space means, that it is possible to interpolate between two latent vectors. Furthermore, every intermediate point in the latent space also corresponds to a pixel pattern that is valid within the distribution of the geometries in the training set.

Therefore, to constrain the inverse design optimization loop to the interpolation regime of the neural network, we have to reformulate the design problem to find an optimum latent vector z instead of an optimum pixel pattern. By optimizing in the WGAN-GP latent space, we can then simply add a regularization term to the inverse design fitness function (we




Figure 2. Average fidelity (a) and MSE values (b, c) through the process of iterative improvement. Values presented are averages for resimulation of 1000 pixel patterns produced to implement random unitary targets. Each improvement loop contained 2000 patterns which are in turn doubled by using y -symmetry.

use mean square error between target and predicted transmission matrix). Since the statistical properties of the WGAN-GP's latent space are known (assuming successful training of the latter), this regularization term has to penalize solutions with latent values far from the mean of the normal distribution. Practically, we add a simple rectified linear unit (ReLU) constraint term to the design fitness function $\sum_i \text{ReLU}(|z_i| - 2)$, that penalizes values of z_i outside of a 2σ range (see Figure 1c).

Inverse Design Loop. A full schematic of our inverse design pipeline is shown in Figure 1c. A set of several latent vectors is randomly initialized. The WGAN generator (blue) predicts the corresponding geometries, of which the complex transmission matrices are predicted by the forward network (green). The latter are compared to the design target transmission using the mean square error as a metric. The total fitness comprises an additional constraint term to restrict the optimization to a 2σ region of the WGAN latent space, which should roughly correspond to the interpolation regime of the neural networks. The gradients of the fitness are calculated using tensorflow's automatic differentiation capability. The initial latent vectors are updated according to the fitness gradients and the process is repeated. After convergence, the best solution is kept as final design.

This process is able to faithfully create pixel patterns that result in distinctly different optical outputs, depending on which input mode is injected to the device. The varFDTD simulated result of the design target shown in Figure 1c is depicted in Figure 1b, showing the successful implementation of our example target, the antidiagonal exchange matrix, with equal phase for light injected in each of the input ports. This method is also capable of addressing the phase of each matrix element, as will be demonstrated in the following sections. Transmission matrices are depicted as colorful sets of 3×3 blocks, where the phase and amplitude of each complex matrix element are mapped respectively on the hue and saturation of an HSV color space.³⁷

FURTHER OPTIMIZATION OF DESIGN PERFORMANCE

Inverse Design Performance Metrics. Inverse design performance is assessed using varFDTD simulation of the optimized patterns on a test-set of 1000 Haar random unitary transmission matrix targets. A useful metric to carry out this analysis is the amplitude fidelity. The fidelity compares both the target and resimulated complex transmission matrices, with a perfect match between the two yielding a fidelity of 1. We

calculate the amplitude fidelity using the following formula taken from literature,⁷³ $F = 1/N[\text{Tr}(|U^*| \cdot |U_{\text{sim}}|)]$ where U^* represents the conjugate transpose of the target complex transmission matrix, U_{sim} is the resimulated result, and N refers to the number of modes in the given system, for our device model, $N = 3$.

Amplitude fidelity as a metric is generally insensitive to the phase agreement between the two matrices under study, as demonstrated in Supporting Information, Figure S8. While intensity based applications such as the development of optical power splitters and routers do not rely on exact phase reconstruction, and some experimental realizations are able to use external phase shifters at the input/output waveguides, this is not always a suitable approach. In applications such as boson sampling or optical computing which require precise control over the relative phases between individual transmission matrix elements, the addition of external phase shifters lacks sufficient degrees of freedom to control all available phases in the matrix, therefore necessitating an optimized pattern to work well across both the real and imaginary components of the electromagnetic field. In such applications, it is instead useful to analyze the mean squared error (MSE) function between target and resimulated phases to assess network performance. The MSE values presented hereafter are calculated across all nine matrix elements. The phase MSE was furthermore taken using the closest phase difference modulo 2π and was normalized by dividing by a factor π^2 . In our analysis we consider mainly the amplitude fidelity as a measure of performance for consistency with other works in the field, but will refer to the other metrics where these are most relevant.

Data Set Optimization. After training both networks on the initial data set (see above), the average amplitude fidelity achieved is around 0.86. While this means in general a good agreement between design target and optimized solution is reached, it also means that the inverse designed patterns still deviate distinctly from the expected results. We therefore invest in optimization of the training data before assessing the performance of the inverse design approach.

Iterative Data Improvement. As a first measure for data set optimization, we iteratively extend the training data set using the results from the inverse design process itself, to improve the design fidelity. Essentially, the idea is to let the neural network predictor learn from its own mistakes.^{36,74} Using the inverse design pipeline, we predict pixel patterns for 2000 random unitary targets. The transmission matrices of the designed patterns are simulated by varFDTD, and these results

are appended to the initial training data set. Subsequently, the models (both forward predictor and WGAN-GP) are retrained using the now extended data set. We repeat this process three times, and benchmark the inverse design quality using 1000 Haar random unitary targets after each iteration. After three iterations, we observe a notable improvement in inverse design performance, with an average amplitude fidelity now of 0.92, as shown in Figure 2. This performance is comparable to fidelities obtained by similarly available reconfigurable technologies,^{73,75,76} which achieve average amplitude fidelities in excess of 0.9.

Expansion with Random Pixel Patterns. The CWGs exhibit a well-defined small number of modes and randomly generated patterns result in high device throughput, as opposed to for example multimode interferometers (MMIs) where the number of modes inside the device is much higher and therefore high-throughput solutions are more sparse.³⁶ This property in principle allows us to rapidly add patterns to the training data set without the need to optimize for device throughput. It may therefore seem appealing to simply append as many random patterns as possible to the training set. However, in our studies we have found a more methodical approach is necessary, as is shown in more detail in Supporting Information, Figure S6, and simply adding more random data does not improve the network performance.

The reason that purely random data sets are found to be ineffective is attributed to the collective action of the pixel patterns. Within the large space of possible solutions, there are many that are not particularly effective in implementing a transmission matrix, as pixels do not sufficiently act in concert or even counteract each other. By including randomly generated data we modify the statistical distribution, which starts containing more and more of these “counterproductive” patterns. Part of the forward model’s capacity is then used to fit these useless noise cases which are not beneficial for the later inverse design task. The WGAN latent space on the other hand will become less efficient in restricting the optimization to useful geometries, since it is also trained on the same data now containing totally random patterns, which are therefore part of its latent space as well. Consequently it is foreseeable that the addition of too many random patterns will result in a decrease of overall inverse design accuracy, which is demonstrated in the Supporting Information, Figures S6 and S7.

Adding a small amount of totally random pixel patterns slightly improves the inverse design capacity, yielding decreased validation losses, as well as an increase in the average fidelity from 0.921 to 0.935, as further shown in Figure 2 (cases “with noise”). This trend is consistent with other work where a small amount of randomness was found to improve performance.⁷⁷ We attribute this improvement to a better prediction of edge case geometries. As pattern optimization happens inside the WGAN latent space, the inverse design is restricted to the statistical distribution of the initial training data. Since the initial data has a bias toward low pixel numbers it may actually be beneficial to also add totally random pixel patterns in order to diversify the training set and make the forward network better deal with patterns at the edge of its latent space

Pattern Characteristics. Study and comparison between the characteristics of pixel patterns from both the training data and inverse designed examples allow us to identify areas where network predictions begin to diverge from geometries used for training. This can therefore serve as a guide for fine-tuning and

development of the network, as well as in training data generation for future geometries or iterative improvement loops. The bottom panel in Figure 3 shows histograms of the




Figure 3. (a) Distribution of pixel numbers contained in patterns from the initial “brute force” training data set (orange) and final training data set after iterative improvement and 2.5% random patterns (green). (b) Pixel histogram of predicted patterns for 1000 random unitary targets. (c) Corresponding average fidelities for each bin of the predicted patterns, with error bars representing the standard deviation from the mean. Dashed lines: average values of 28 pixels for initial training data and 41 pixels for final training data.

distribution of active pixels in the initial data sets used for training (orange) and the augmented data sets after three iteration loops (green). The pixel histogram for patterns predicted by the network for a validation data set of 1000 target transmission matrices are shown in the middle panel (blue). The corresponding average and standard deviation of the fidelity for the predicted patterns is presented in the top panel of the figure. As the training data is generated in an iterative manner and intermediate steps are stored, there is an strong emphasis in the training data set for patterns contain a low number of pixels. Across the entire database used for training there is an average number of 41 pixels/pattern.

In comparison, the average number of pixels in predicted patterns is 104 and patterns with up to 225 active pixels are found. Implementing the target transmission matrices requires predominantly geometries containing more pixels than the average number from the training data. We can observe that as the number of pixels increases above 50, the standard deviation of the fidelity increases quite significantly. An increase in network error rate will lead to a greater spread of results and thus a larger standard deviation. This trend is attributed to the forward network performing less accurately on geometries that are underrepresented in its training data, which could explain some of the errors in the inverse design. More discussion on the influence of forward network accuracy on fidelity will be given below.




Figure 4. Performance of predicted patterns for the implementation of all available permutation matrices of the $n = 3$ waveguide array. 200 random phase shifts are applied to each matrix element, allowing for in depth sampling of the available phase space to determine a maximum fidelity for phase insensitive applications. Average fidelities of around 0.9 are achieved for all targets, with peak fidelity values >0.97 for all targets.

RESULTS ON TRANSMISSION MATRIX DESIGN

After the iterative data set expansion, the addition of random patterns, and the training of the network on this final data set, we proceed to carry out a benchmark of our approach for the inverse design of pixel patterns for a variety of coupled waveguide transmission matrices. These can be broadly split into two groups: permutation matrices, which are a class of orthogonal matrices where individual input and output ports are connected without port mixing, and unitary matrices, providing the most general input–output relationship. In the following we discuss these cases in more detail.

Permutation Matrices. The first class of matrices under study contains those which guide light in a one-to-one fashion, whereby light injection into each input waveguide results in transmission through only one unique output. For a transmission matrix with $n \times n$ elements, there will be $n!$ unique permutation matrices, each of which for our device can be seen in Figure 4. An example phase shifted matrix is presented for each permutation target as well as its associated varFDTD

simulation of the near field electric field intensity. These matrices are a useful test to ensure that the network has generalized to a point where it is able to predict patterns for targets which lie on the very edge of the training geometry space (those confining the output electric field fully to only one port), and to check the ability of inverse designed patterns to produce distinctly different optical outputs for each input mode. In our design challenge, we chose to set the phase of each nonzero matrix element as a free parameter which can be chosen to obtain the best solution in amplitude. This approach appears useful since (i) in many use cases, the phase response of an optical switch is not critical and (ii) if necessary, an optical phase shifter at each of the outputs is sufficient to rebalance the output phases in the device.

Results shown in Figure 4 demonstrate accurate reproduction of both the amplitude and phase distribution for a selection of matrices, demonstrating that a well performing solution can be achieved for all of the orthogonal permutation matrices in the 3×3 group. Here solutions were selected on




Figure 5. (a) Distribution of fidelities for the reproduction of 1000 random unitary targets. Highlighted are instances which fall outside 1 and 2σ of the average MSE between forward network prediction and resimulation, indicating a failed forward network prediction. (b) and (c) show the distribution of MSE values between target and resimulation for these same targets. (d) Four example unitary targets sampled from across the achieved fidelity distribution. We present the target, forward network prediction of the inverse designed pixel pattern, and a resimulated transmission matrix. To the right-hand side we plot the respective near field electric field intensity map for a device programmed with the predicted pattern for each target matrix.

the basis of combining fidelity with a low phase MSE. Pixel patterns as well as local field distributions for each of the input ports are presented for completeness for each of the devices, providing some insight in the internal structure of these solutions. For example, for the identity matrix at the top, we see that the pattern of perturbations is predominantly placed in the middle of the device along its length. This pattern transforms the original device, which had no self-imaging capability, to produce a highly symmetric self-imaging of each of the three inputs onto the output plane.

Antidiagonal and more complex permutations are seen in the other five diagrams, where the common denominator is that the field profiles typically show only a few bounces between input and output, indicative of the coupling length of the CWG of around $10\ \mu\text{m}$ and the fact that the modal basis is very small. Weak perturbations therefore tend to couple modes in an adiabatic way, without giving rise to strong scattering events or reflections.³⁴

A sample of fidelities across the available phase space for each permutation matrix is presented in the right column in Figure 4, in which 200 random phase shifts are applied to each element in the target matrix. Average fidelities are achieved of around 90% across the phase space sample for all 6 target matrices. When output phases can be freely adjusted to maximize the fidelity, the phase space sampling technique permits significant improvements to the fidelity, resulting in fidelities in excess of 0.97 for all targets.

General Unitary Transmission Matrices. Next to the design of orthogonal permutation matrix targets, which pose a useful design task to ensure that we have precise control over every individual matrix element, in many real world use cases such as optical computing and quantum information processing, mixtures of inputs are required. Unitary matrices represent a particularly useful class of targets representing the most general operations available in the multiport system. These targets must satisfy the condition $UU^* = I$, where U^* represents the conjugate transpose of the target matrix and I is the identity matrix, corresponding to a lossless device used for target matrices.

Figure 5a shows the fidelity distribution for network predicted patterns from a set of 1000 random unitary targets drawn from the Haar distribution. We find that the neural network is able to reach design targets with good fidelity of 0.935 ± 0.038 . This is reflected in the convergence to low values below 0.05 in the amplitude MSE. The distribution of MSE values for phase shows that the prediction of correct phase distributions is more challenging. A subselection of four matrix targets from across the fidelity distribution are presented in Figure 5b, labeled as (i)–(iv), as indicated by the colored dots in Figure 5b, showing the target matrix, the forward network prediction and varFDTD resimulation result. Pixel patterns and near field maps are present the corresponding microscopic configurations of the four patterned devices under study. It can be seen that for all four examples, the forward network prediction agrees well with the




Figure 6. (a) Polar plot comparing each matrix element in the reproduction of the 3×3 Hadamard and Fourier matrix (shown in b). Data points are color coded to their associated target value, and black lines indicate the distance on the complex plane between forward network predictions and their resimulated matrix values. (c) Phase sweep is shown again, in which a global phase shift is applied to the target matrix, retaining the same intensity distribution resulting in a maximum achievable fidelity in excess of 0.99.

target matrix. Main differences can be seen between the network prediction and the varFDTD result, suggesting that an important factor in the fidelity may be the accuracy of the forward network.

To verify the influence of the forward network error, we calculate the MSE between forward network prediction and varFDTD calculation. The fidelity distributions are obtained for parts of the data set where the MSE lies above one and two standard deviations (labeled as 1σ and 2σ), given by respectively the orange and green histograms in Figure 5a. This analysis confirms what is seen in the examples, namely, that the tail of lower fidelities is correlated with a poor accuracy of the forward network prediction.

The examples furthermore indicate that phase and amplitude errors are not related. Results with low phase error can be found for a low fidelity and vice versa. Figure 5d(iii), for example, shows visually a good match to the target unitary colors due to low phase error; however, the final amplitude fidelity remains below average at 0.9, owing to a disagreement between amplitude values. The near field maps and pixel patterns of random unitaries are not easily understood through intuition, while an underlying phenomenology would be of interest this goes beyond the scope of our study. We do observe the emergence of larger section of connected pixels forming lines, which may be a strategy of the network to achieve large phase shifts in certain matrix elements. Some more discussion on the phase structure of the CWGs is presented in our Discussion section.

Hadamard and Fourier Matrix. Complex Hadamard matrices play an important role in quantum information theory. They have been used to tackle a number of problems including the development of spin models⁷⁸ and analogue quantum simulators.⁷⁹ They have also helped establish mathematical frameworks to construct bases of unitary operators and maximally entangled states. Fourier matrices can be used to apply a discrete Fourier transform to a signal through matrix multiplication, making them of particular interest in optical signal processing and computation. For a 3×3 matrix, all complex Hadamard matrices are equivalent to the Fourier matrix, F_3 :⁸⁰

$$F_3 = \begin{bmatrix} 1 & 1 & 1 \\ 1 & \omega & \omega^2 \\ 1 & \omega^2 & \omega \end{bmatrix}$$

where $\omega = \exp(2\pi i/3)$. It is of interest to check the existence of a solution for this specific matrix operator. As F_3 is a unique matrix, we again give the network some more flexibility in the boundary conditions by allowing a single global phase factor within all matrix elements. A global phase factor can easily be factorized out and compensated in an optical system. It makes sense to include this degree of freedom in the matrix to find the best working point of the CWG under study, taking into account both amplitude and phase MSE. Figure 6a,b shows the forward network prediction and resimulated result when attempting to implement the F_3 matrix. An amplitude fidelity of 0.94 is achieved for the selected solution, with a phase MSE of 0.035 and amplitude MSE of 0.0145. The polar plot in Figure 6a is used as a graphical depiction of the accuracy for each individual matrix element, data points are color coded corresponding with the matrix element they represent, with black lines between associated network predictions and resimulated points representing their separation within the complex plane.

The full histogram of fidelities for a range of global phase values between $-\pi$ and π is shown in Figure 6c and allows to identify an average fidelity across this phase sample of 0.93, while we again are able to retrieve a peak fidelity of >0.99 , but at a higher phase MSE.

A question of considerable importance is the stability of the obtained solutions against environmental variations, and particularly changes in temperature. We used the specific design of the Fourier matrix to explore this point in more depth in Supporting Information, Figure S11. Here, we have taken into consideration the thermo-optic coefficient of silicon⁸¹ as the main contribution for simulating the matrix for a range of temperatures around room temperature. It is found that, while the complex matrix elements show a significant temperature effect, most of this can be extracted as a global thermo-optic phase shift acting equally on all matrix elements, with a much smaller effect on the fidelity and on the relative phases between the matrix elements. We therefore conclude that the matrix design is robust against temperature variations over a range of tens of Kelvins.

DISCUSSION

Our work shown here presents coupled waveguide arrays as an alternative to other commonly used reconfigurable platforms such as interferometer meshes or multimode based devices for

programming arbitrary unitary operators. Approaches for using such devices are well established, with practical realizations demonstrated across a number of fields as discussed in the preamble to this work. Interferometer meshes in particular provide a high degree of control, allowing for the introduction of individual phase delays at single unit cells across the mesh for highly accurate matrix reproduction. There exist several mathematical models^{82,83} which allow for the calculation of where these phase shifts must be applied for arbitrary matrix decomposition. Despite promising results, small insertion losses at each interferometer will sum together, leading to non-negligible losses across the whole device. Furthermore, when targeting large matrices, the total device footprint becomes large when compared to alternatives. Adding to this, many approaches use thermo-optic modulators to introduce the required phase shifts, which add further to this size, introduce electronic overhead and regulating circuits, and produce thermal instability across the chip.

Additional degrees of freedom may be offered also to CWGs by introduction of external phase shifters. Already in Figure 6 we have considered the freedom of adding a global phase offset to all matrix elements to find an optimum solution in terms of both phase MSE and fidelity. A global phase shift is an operation which does not change the internal phases between the matrix elements and can be easily achieved externally. Also in orthogonal permutation matrices many elements are zero and phase shifters at each of the outputs can readily control the relative phases for all nonzero matrix elements. This extra design freedom offered by choosing an arbitrary phase has allowed some “cherry picking” of solutions in parameter space as shown in 4. For the full unitary matrix of an $N \times N$ multiport device there exist N^2 transmission matrix elements but only $2N$ input/output waveguides at which we may place phase shifters, therefore we lack sufficient degrees of freedom to tune the phase of every matrix element. In Figure 5 we restricted ourselves to zero freedom in phase to show the ability and limitations of the method in achieving random unitary matrices, however in applications one may make use of the $2N$ external phase shifters to also reduce some of the challenges in unitary matrix design.

As an alternative approach to CWGs, multimode structures such as MMIs and devices which are built around multimode waveguides are the most compact technology discussed here. Extra degrees of freedom introduced by the inclusion of different optical modes allow for highly compact matrix decomposition, however, this is not without its drawbacks. Coupling out of the device becomes challenging as the MMI has many more internal degrees of freedom than the input and output modes, therefore any modal-mismatch can introduce significant losses, often requiring some signal preprocessing through techniques such as beam shaping. The need for accurate coupling between modes also now imposes a requirement for high degrees of spatial accuracy in the programming of such a device. Coupled waveguide arrays strike a midpoint in device footprint, necessitating only slightly more space in the direction of light propagation to allow a round trip between outermost waveguides if aiming for the smallest device possible. In the plane perpendicular to the light injection axis their size remains comparable, if not smaller than MMI type devices. This compact form-factor is significantly smaller than comparable interferometer meshes, and remains within a consistent modal basis, introducing no losses from

coupling between modes at the output waveguides making it an attractive option for space-limited applications.

While in general the network performance in our work is strong, there are some examples where the resimulated results are not a good match to the target matrix. By comparing the forward network predictions to resimulated results we may check if forward network predictions remain valid. Using a simple MSE discriminator, instances which fall outside one standard deviation from the mean can be identified and rejected, highlighting cases where the forward network has failed to accurately predict the transmission matrix for a given pattern. The fidelity distribution for cases with resimulation error above one and two standard deviations (above 1 and 2σ) are highlighted in orange and green in the histogram of Figure 5. For these designs, the network appears to have strayed from the interpolation regime of the forward network, resulting in large errors in the predicted transmission matrix. Rather unsurprisingly, almost every fidelity below 0.85 is a result of a failed network prediction. Excluding these from our statistical analysis the average fidelity now raises to 0.94 ± 0.029 . In the majority of cases, however, it appears that the optimization of the geometry in the WGAN's latent space and the addition of the latent constraint successfully limits designs to the surrogate network's validity region. Iterative data improvement may also be continued beyond the three cycles in this study, which apart from making the network learn from its own mistakes allows for augmenting the training data toward larger pixel numbers. The performance gain has to be traded off against the effort needed as the additional improvement for each cycle is expected to saturate.

The network is trained on the full complex field, therefore one might expect similar performance for phase and amplitude, study of our respective MSE values shows this is, however, not the case. One hypothesis for the cause of this discrepancy originates in how our brute force optimization targets are generated. During the data generation stage, the phase of light at the outputs is unconstrained and given no optimization weighting, we assess the success of each pixel iteration solely on the intensity at the output ports. Intensity optimization targets are randomly selected from a uniform distribution, resulting in the retrieval of a normal distribution of possible amplitude values centered around 0.33 as we expect for a 3-waveguide system. Conversely, analysis of the phase distribution for the same data (shown in Supporting Information, Figure S10) shows distinctly different behavior. Unsurprisingly we record a peak for phase values matching to those of an unperturbed device, however, because the optimization is not pushed to suggest pixel patterns with “extreme” phase delays approaching 2π , the data set becomes highly biased toward smaller phase shifts. The final database of pixel patterns represent the available intensity space well, but may fall short when describing the phase space especially for longer phase delays. Consequently it follows that the network's ability to accurately predict patterns to implement such phase delays will be inhibited, explaining the differing MSE results for phase and amplitude we observe. The addition of an additional phase target into the brute force algorithm has little effect on the final phase distribution. Each matrix element may only be delayed from its initial unperturbed phase, meaning regardless of how targets are defined, the data set will always contain more examples of small phase shifts as these originate from the start of the optimization where pixels are more readily accepted. Therefore, if phase accuracy is essential a training data set

should be created using an adapted approach to that outlined here, for example one composed solely from random patterns.

The major limiting factor in the final generation of valid pixel patterns for user-defined matrix targets remains the accuracy with which the forward network can make predictions. Assessment of this will give some insight as to the expected upper limit on network performance when presented with an external design target. [Supporting Information, Figure S9](#) shows the distribution of fidelities from forward network predictions on the validation data set. This represents a "best case" scenario, as these patterns are generated in precisely the same way as the training data. Across the 6564 validation patterns we record an average fidelity of 0.92 ± 0.04 , indicating that for our current predictor model, the achieved average amplitude fidelity of 0.94 constitutes what is likely a near maximal achievable value, which is limited by the network performance rather than the physical CWG system. More conventional topology optimizations starting from the end point of the neural-adjoint inverse design may be of interest to further converge solutions, which could be a topic of future work.

Finally we point out a paradox between first the role of the WGAN-GP in constraining the pattern predictions within the interpolation regime of the forward predictor network, and second the large increase in number of pixels for the design solutions versus the training data set seen in [Figure 3](#). We have observed similar behavior in previous work on patterned MMIs³⁶ in which the pixels contained in deep learning predictions differ significantly from those seen in the networks training data set. Many pixels are required to optimize the pattern for the entire complex transmission matrix at once, especially when using thin films of PCM which introduce modest effective index variations at each perturbation, although this may be reduced by increasing the thickness of phase change material.⁸⁴ The training data set is highly biased toward smaller pixel numbers because we save patterns iteratively starting from zero pixels. Importantly, the majority of the training data set is directed to only optimize one row of the matrix at a time, i.e. for one input port, and not the full matrix itself. To achieve design of the full matrix, the optimization algorithm has to combine knowledge on patterns for the three different input ports. Because the latent vectors exist inside a high dimensional space it is difficult to gain an intuition for which patterns will lie where within the space, however, it is likely that patterns with more pixels present will lie further from the center. The general agreement between simulation and forward network predictions in the majority of cases indicate that the WGAN-GP has effectively confined predictions to those within the validity region of the forward net. Future work could explore how to further optimize this interplay at the boundary of latent space in order to further improve performance.

CONCLUSIONS

In conclusion, our study addresses the potential of low-loss optical phase change materials combined with coupled waveguide arrays as a promising and exciting avenue for the production of next-generation reconfigurable technologies across fields such as quantum simulation, photonic computing and optical data processing. In comparison with other approaches such as integrated circuits based around interferometer meshes or MMIs, coupled waveguide arrays can offer ultralow-loss reversible modulation requiring no active

regulation within a highly compact device footprint. The development of a robust inverse design pipeline using neural network surrogate models allows for rapid, near-real-time prediction of the complex pixel patterns, required to implement a wide range of transmission matrices within a single device model. In our work, the introduction of a Wasserstein generative adversarial network provides a crucial constraint on the gradient based optimization, limiting predicted geometries to those within the interpolation region of the forward surrogate model. Network performance is enhanced by augmenting and expanding the training data set, initially through a process of iterative improvement, and in a final step, with the introduction of a small percentage of noisy random data. Although training of the networks is time-consuming, taking a few hours with standard consumer level computer hardware, the entire inverse design process remains significantly faster than alternatives such as topology optimizations, allowing predictions on a millisecond time scale once trained.

Presented results demonstrate a high level of control over both intensity and phase of individual matrix elements. While some extreme phase relations may suggest the need to expand the waveguide geometries to longer device, therefore allowing multiple vertical passes of the light, generally performance is strong. Average fidelities of 0.935 ± 0.04 are reported, demonstrating comparable performance to other state of the art, commercially available reconfigurable photonic technologies and further reinforcing the validity of this approach. The fidelity obtained in our study is currently limited by the neural network performance rather than the coupled waveguide system itself and further improvement may be possible in future work. The neural-adjoint design platform introduced is highly versatile and requires minimal overheads to allow functionality across different devices and geometries, showing promise for integration with a range of future optical technologies although as devices grow in scale the available perturbation pattern geometry space grows exponentially, requiring extra simulation time to create a representative data set. The exact point at which this becomes prohibitively long will be application specific and dependent on how frequently new patterns must be designed but will at some point become a limitation of our approach.

METHODS

Device Simulations. Coupled waveguide devices are simulated using the finite difference time domain (FDTD) software Lumerical, specifically we use the Variational FDTD (varFDTD) engine, MODE. This engine collapses a 3D device structure into an effective 2D simulation by probing the effective index experienced in slices across the simulation mesh. The resultant simulation achieves comparable accuracy to a full 3D simulation, however requires significantly less computational resources. As this method accounts for the 3D structure of the device model it is commonly referred to as a 2.5D simulation. The varFDTD approach allows for a dramatic reduction in the simulation time required to create a representative data set of pixel patterns and serves as an approximation of a full-vectorial 3D modal evolution of the light within our devices. [Supporting Information, Figure S10](#) shows the agreement we observe between the two simulation regimes for a patterned waveguide device. In the design, waveguides are obtained by etching 120 nm deep into a 220 nm thick silicon overlayer, which lies on top of a 2 μm thick

buried oxide layer. The rib waveguides are 500 nm wide, allowing only propagation of the fundamental mode at the considered vacuum wavelength of 1550 nm. The distance between waveguides in the CWG section is 250 nm over a 50 μm long coupled region. To avoid any cross-talk outside of this coupling region, waveguides fan out to a spacing of 1 μm at both input and output sides. The waveguides are covered with a 30 nm thick Sb_2Se_3 layer to maintain their single mode performance, displaying an index contrast between crystalline and amorphous states of $\Delta n_{\text{Sb}_2\text{Se}_3} = 0.77$. Finally, the structures are capped with a semi-infinite cladding layer of SiO_2 . The complex dielectric function for Sb_2Se_3 was taken from ellipsometry measurements.²⁰ Simulations on ideal waveguides do not take into account additional in-plane propagation losses due to scattering from polycrystalline grain boundaries and surface roughness, which play a predominant role in crystalline Sb_2Se_3 films of increasing thickness.⁸⁴

Training Data Generation Using Iterative Optimization. The brute force iterative optimization process starts by randomly choosing a specific input channel, as well as a random splitting ratio for the output channels. Then, the state of a single, random pixel is switched and the device transmission matrix is evaluated through varFDTD simulations. If, compared to the previous device geometry, the result is closer to the target splitting ratio, the pixel is retained and further simulations are run for each input channel in order to record the full transmission matrix. If not, the pixel is reverted to its initial state, another random pixel is selected, and the process is repeated. After a specified number of iterative cycles the process starts over with a blank device and a new optimization target. A more detailed description of this process as well as a specific example can be found in the [Supporting Information, Figure S5](#). After generation of the data set, the horizontal plane of symmetry found in the unperturbed waveguide array was furthermore exploited so as to double the size of this training set.

Deep Learning Neural Network. The forward predictor used is a ResNeXt encoder-decoder convolutional neural network,⁶⁹ a detailed schematic of this forward network model can be found in the [Supporting Information, Figure S1](#). Training of the network takes about 1 h using a Nvidia RTX3070 GPU. To avoid convergence of solutions to the extrapolation domain, a Wasserstein generative adversarial network with gradient penalty⁷² (WGAN-GP) was used to develop a learned reparametrization of the pixel patterns geometry from the training data set. Details of the WGAN-GP architecture are shown in the [Supporting Information, Figure S2](#). Through the WGAN-GP, the pixel-based representations of the geometries that form a nonconvex set in the original parametrization (i.e., interpolation leads to nonphysical patterns with gray scale pixels) are mapped into a convex set of latent representations of the same geometries. Any interpolation between these latent vectors should still represent a valid pattern, which can be generated from its latent vector using the WGAN-GP generator network. This also means that due to the training procedure using normally distributed random sampling, all geometries that are valid interpolations of the data set samples lie within said normal distribution (with known $\mu_z = 0$, $\sigma_z = 1$). A rectified linear unit (ReLU) constraint term to the design fitness function $\sum_i \text{ReLU}(|z_i| - 2)$, that penalizes values of z_i outside of a 2σ range.

Gradient Descent. Gradient-based optimizations such as used here are prone to become stuck in local minima and

therefore require good initial guesses to reach a globally optimum solution. To ensure that we find a close to ideal solution, we run a large number of concurrent optimizations for each design target (100 in our case). Eventually, the best ranked solution is kept. We note that due to the highly optimized GPU-based parallelization of modern deep learning frameworks such as tensorflow used in this work, the concurrent optimization of many targets is very efficient and fast.

■ ASSOCIATED CONTENT

Data Availability Statement

Supporting data used in this work is openly available from the University of Southampton repository at doi.org/10.5258/SOTON/D3348.

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acsp Photonics.4c02081>.

Technical details regarding the structure of the presented neural network, training data generation process, and supporting data for the analysis of results (PDF)

■ AUTHOR INFORMATION

Corresponding Authors

Thomas W. Radford – School of Physics and Astronomy, University of Southampton, Southampton SO17 1BJ, United Kingdom; orcid.org/0000-0002-3410-3206; Email: t.radford@soton.ac.uk

Otto L. Muskens – School of Physics and Astronomy, University of Southampton, Southampton SO17 1BJ, United Kingdom; orcid.org/0000-0003-0693-5504; Phone: +44 (0)23 80593911; Email: o.muskens@soton.ac.uk

Authors

Peter R. Wiecha – LAAS, Université de Toulouse, CNRS, 31031 Toulouse, France; orcid.org/0000-0002-4571-0116

Alberto Politi – School of Physics and Astronomy, University of Southampton, Southampton SO17 1BJ, United Kingdom; orcid.org/0000-0002-3668-9474

Ioannis Zeimpekis – School of Electronics and Computer Science and Optoelectronics Research Centre, University of Southampton, Southampton SO17 1BJ, United Kingdom

Complete contact information is available at:

<https://pubs.acs.org/doi/10.1021/acsp Photonics.4c02081>

Funding

This work was supported financially by EPSRC through a Ph.D. studentship for T.R. from the Doctoral Training Partnership EP/T517859/1, Project Reference 2611785. A.P. acknowledges support from EPSRC through Grant EP/X016749/1.

Notes

The authors declare no competing financial interest.

■ REFERENCES

- (1) Zhou, H.; Dong, J.; Cheng, J.; Dong, W.; Huang, C.; Shen, Y.; Zhang, Q.; Gu, M.; Qian, C.; Chen, H.; et al. Photonic matrix multiplication lights up photonic accelerator and beyond. *Light: Science & Applications* **2022**, *11*, 30.

- (2) Giamougiannis, G.; Tsakyridis, A.; Moralis-Pegios, M.; Mourgas-Alexandris, G.; Totovic, A. R.; Dabos, G.; Kirtas, M.; Passalis, N.; Tefas, A.; Kalavrouziotis, D.; et al. Neuromorphic silicon photonics with 50 GHz tiled matrix multiplication for deep-learning applications. *Advanced Photonics* **2023**, *5*, 016004.
- (3) Aspuru-Guzik, A.; Walther, P. Photonic quantum simulators. *Nat. Phys.* **2012**, *8*, 285–291.
- (4) Wang, J.; Sciarino, F.; Laing, A.; Thompson, M. G. Integrated photonic quantum technologies. *Nat. Photonics* **2020**, *14*, 273–284.
- (5) Choutagunta, K.; Roberts, I.; Miller, D. A.; Kahn, J. M. Adapting Mach–Zehnder mesh equalizers in direct-detection mode-division-multiplexed links. *Journal of Lightwave Technology* **2020**, *38*, 723–735.
- (6) Shastri, B. J.; Tait, A. N.; Ferreira de Lima, T.; Pernice, W. H.; Bhaskaran, H.; Wright, C. D.; Prucnal, P. R. Photonics for artificial intelligence and neuromorphic computing. *Nat. Photonics* **2021**, *15*, 102–114.
- (7) Bogaerts, W.; Rahim, A. Programmable Photonics: An Opportunity for an Accessible Large-Volume PIC Ecosystem. *IEEE J. Sel. Top. Quantum Electron.* **2020**, *26*, 1–17.
- (8) Bogaerts, W.; Pérez, D.; Capmany, J.; Miller, D. A. B.; Poon, J.; Englund, D.; Morichetti, F.; Melloni, A. Programmable photonic circuits. *Nature* **2020**, *586*, 207–216.
- (9) Chen, X.; Milosevic, M. M.; Thomson, D. J.; Khokhar, A. Z.; Franz, Y.; Runge, A. F.; Mailis, S.; Peacock, A. C.; Reed, G. T. Post-fabrication phase trimming of Mach–Zehnder interferometers by laser annealing of germanium implanted waveguides. *Photonics Research* **2017**, *5*, 578–582.
- (10) Zheng, J.; Khanolkar, A.; Xu, P.; Colburn, S.; Deshmukh, S.; Myers, J.; Frantz, J.; Pop, E.; Hendrickson, J.; Doyle, J.; et al. GST-on-silicon hybrid nanophotonic integrated circuits: a non-volatile quasi-continuously reprogrammable platform. *Optical Materials Express* **2018**, *8*, 1551–1561.
- (11) Staffoli, E.; Mancinelli, M.; Bettotti, P.; Pavesi, L. Equalization of a 10 Gbps IMDD signal by a small silicon photonics time delayed neural network. *Photonics Research* **2023**, *11*, 878–886.
- (12) Shen, Y.; Harris, N. C.; Skirlo, S.; Prabhu, M.; Baehr-Jones, T.; Hochberg, M.; Sun, X.; Zhao, S.; Larochelle, H.; Englund, D.; et al. Deep learning with coherent nanophotonic circuits. *Nat. Photonics* **2017**, *11*, 441–446.
- (13) Pérez, D.; Gasulla, I.; Crudgington, L.; Thomson, D. J.; Khokhar, A. Z.; Li, K.; Cao, W.; Mashanovich, G. Z.; Capmany, J. Multipurpose silicon photonics signal processor core. *Nat. Commun.* **2017**, *8*, 636.
- (14) Ashtiani, F.; Geers, A. J.; Aflatouni, F. An on-chip photonic deep neural network for image classification. *Nature* **2022**, *606*, 501–506.
- (15) Wu, D.; Wu, Y.; Wang, Y.; An, J.; Hu, X. Reconfigurable optical add-drop multiplexer based on thermally tunable micro-ring resonators. *Opt. Commun.* **2016**, *367*, 44–49.
- (16) Atabaki, A. H.; Eftekhari, A. A.; Yegnanarayanan, S.; Adibi, A. Sub-100-ns thermal reconfiguration of silicon photonic devices. *Opt. Express* **2013**, *21*, 15706–15718.
- (17) Errando-Herranz, C.; Niklaus, F.; Stemme, G.; Gylfason, K. B. A low-power MEMS tunable photonic ring resonator for reconfigurable optical networks. *2015 28th IEEE International Conference on Micro Electro Mechanical Systems (MEMS)*; IEEE, 2015; pp 53–56.
- (18) Oshita, M.; Takahashi, H.; Ajiki, Y.; Kan, T. Reconfigurable surface plasmon resonance photodetector with a MEMS deformable cantilever. *ACS Photonics* **2020**, *7*, 673–679.
- (19) Quack, N.; et al. Integrated silicon photonic MEMS. *Microsyst. Nanoeng.* **2023**, *9*, 27.
- (20) Delaney, M.; Zeimpekis, I.; Lawson, D.; Hewak, D. W.; Muskens, O. L. A new family of ultralow loss reversible phase-change materials for photonic integrated circuits: Sb₂S₃ and Sb₂Se₃. *Adv. Funct. Mater.* **2020**, *30*, 2002447.
- (21) Miller, K. J.; Haglund, R. F.; Weiss, S. M. Optical phase change materials in integrated silicon photonic devices: review. *Optical Materials Express* **2018**, *8* (8), 2415–2429.
- (22) Chen, R.; Fang, Z.; Miller, F.; Rarick, H.; Fröch, J. E.; Majumdar, A. Opportunities and Challenges for Large-Scale Phase-Change Material Integrated Electro-Photonics. *ACS Photonics* **2022**, *9*, 3181–3195.
- (23) Wu, C.; Deng, H.; Huang, Y.-S.; Yu, H.; Takeuchi, I.; Ríos Ocampo, C. A.; Li, M. Freeform direct-write and rewritable photonic integrated circuits in phase-change thin films. *Science Advances* **2024**, *10*, No. eadk1361.
- (24) Prabhathan, P.; et al. Roadmap for phase change materials in photonics and beyond. *iScience* **2023**, *26*, 107946.
- (25) Tripathi, D.; Vyas, H. S.; Kumar, S.; Panda, S. S.; Hegde, R. Recent developments in Chalcogenide phase change material-based nanophotonics. *Nanotechnology* **2023**, *34*, S02001.
- (26) Wuttig, M.; Bhaskaran, H.; Taubner, T. Phase-change materials for non-volatile photonic applications. *Nat. Photonics* **2017**, *11*, 465–476.
- (27) Delaney, M.; Zeimpekis, I.; Du, H.; Yan, X.; Banakar, M.; Thomson, D. J.; Hewak, D. W.; Muskens, O. L. Nonvolatile programmable silicon photonics using an ultralow-loss Sb₂Se₃ phase change material. *Science Advances* **2021**, *7*, 3500–3516.
- (28) Martin-Monier, L.; Popescu, C. C.; Ranno, L.; Mills, B.; Geiger, S.; Callahan, D.; Moebius, M.; Hu, J. Endurance of chalcogenide optical phase change materials: a review. *Opt. Mater. Express* **2022**, *12*, 2145–2167.
- (29) Lawson, D.; Blundell, S.; Ebert, M.; Muskens, O. L.; Zeimpekis, I. Optical switching beyond a million cycles of low-loss phase change material Sb₂Se₃. *Optical Materials Express* **2024**, *14*, 22–38.
- (30) Alam, M. S.; Laing, R.; Bolatbek, Z.; Heenkenda, R.; Gnawali, R.; Payne, T. E.; Sarangan, A.; Hendrickson, J. R.; Agha, I. Fast Cycling Speed with Multimillion Cycling Endurance of Ultra-Low Loss Phase Change Material (Sb₂Se₃) by Engineered Laser Pulse Irradiation. *Adv. Funct. Mater.* **2024**, *34*, 2310306.
- (31) Miller, D. A. B. Self-configuring universal linear optical component [Invited]. *Photonics Research* **2013**, *1*, 1.
- (32) Liu, V.; Jiao, Y.; Miller, D. A.; Fan, S. Design methodology for compact photonic-crystal-based wavelength division multiplexers. *Optics letters* **2011**, *36*, 591–593.
- (33) Lu, J.; Vučković, J. Nanophotonic computational design. *Opt. Express* **2013**, *21*, 13351–13367.
- (34) Bruck, R.; Vynck, K.; Lalanne, P.; Mills, B.; Thomson, D. J.; Mashanovich, G. Z.; Reed, G. T.; Muskens, O. L. All-Optical Spatial Light Modulator for Reconfigurable Silicon Photonic Circuits. *Optica* **2016**, *3*, 396.
- (35) Frellsen, L. F.; Ding, Y.; Sigmund, O.; Frandsen, L. H. Topology optimized mode multiplexing in silicon-on-insulator photonic wire waveguides. *Opt. Express* **2016**, *24*, 16866–16873.
- (36) Dinsdale, N. J.; Wiecha, P. R.; Delaney, M.; Reynolds, J.; Ebert, M.; Zeimpekis, I.; Thomson, D. J.; Reed, G. T.; Lalanne, P.; Vynck, K.; et al. Deep learning enabled design of complex transmission matrices for universal optical components. *ACS Photonics* **2021**, *8*, 283–295.
- (37) Nikkhah, V.; Pirmoradi, A.; Ashtiani, F.; Edwards, B.; Aflatouni, F.; Engheta, N. Inverse-designed low-index-contrast structures on a silicon photonics platform for vector–matrix multiplication. *Nat. Photonics* **2024**, *18*, 501–508.
- (38) Taballione, C.; Wolterink, T. A.; Lugani, J.; Eckstein, A.; Bell, B. A.; Grootjans, R.; Visscher, I.; Geskus, D.; Roeloffzen, C. G.; Renema, J. J.; et al. 8 × 8 reconfigurable quantum photonic processor based on silicon nitride waveguides. *Opt. Express* **2019**, *27*, 26842–26857.
- (39) Perets, H. B.; Lahini, Y.; Pozzi, F.; Sorel, M.; Morandotti, R.; Silberberg, Y. Realization of Quantum Walks with Negligible Decoherence in Waveguide Lattices. *Phys. Rev. Lett.* **2008**, *100*, 170506.
- (40) Peruzzo, A.; Lobino, M.; Matthews, J. C. F.; Matsuda, N.; Politi, A.; Poulios, K.; Zhou, X.-Q.; Lahini, Y.; Ismail, N.; Wörhoff, K.; Bromberg, Y.; Silberberg, Y.; Thompson, M. G.; O'Brien, J. L. Quantum Walks of Correlated Photons. *Science* **2010**, *329*, 1500–1503.

- (41) Garanovich, I. L.; Longhi, S.; Sukhorukov, A. A.; Kivshar, Y. S. Light propagation and localization in modulated photonic lattices and waveguides. *Phys. Rep.* **2012**, *518*, 1–79.
- (42) Lahini, Y.; Steinbrecher, G. R.; Bookatz, A. D.; Englund, D. Realization of Quantum Walks with Negligible Decoherence in Waveguide Lattices. *njp Quant. Inf.* **2018**, *4*, 2.
- (43) Petrovic, J. Multiport waveguide couplers with periodic energy exchange. *Opt. Lett.* **2015**, *40*, 139–142.
- (44) Petrovic, J.; Kršić, J.; Maluckov, A.; Veerman, J. High-density optical interconnects based on self-imaging in coupled waveguide arrays. *Optics & Laser Technology* **2023**, *163*, 109381.
- (45) Yang, Y.; Chapman, R. J.; Haylock, B.; Lenzini, F.; Joglekar, Y. N.; Lobino, M.; Peruzzo, A. Programmable high-dimensional Hamiltonian in a photonic waveguide array. *Nat. Commun.* **2024**, *15*, 50.
- (46) Bennet, P.; Langevin, D.; Essoual, C.; Khaireh-Walieh, A.; Teytaud, O.; Wiecha, P.; Moreau, A. Illustrated Tutorial on Global Optimization in Nanophotonics. *JOSA B* **2024**, *41*, A126–A145.
- (47) Schneider, P.-I.; Garcia Santiago, X.; Soltwisch, V.; Hammerschmidt, M.; Burger, S.; Rockstuhl, C. Benchmarking Five Global Optimization Approaches for Nano-optical Shape Optimization and Parameter Reconstruction. *ACS Photonics* **2019**, *6*, 2726–2733.
- (48) Bendsoe, M. P. Optimal shape design as a material distribution problem. *Structural optimization* **1989**, *1*, 193–202.
- (49) Jensen, J. S.; Sigmund, O. Topology optimization for nanophotonics. *Laser & Photonics Reviews* **2011**, *5*, 308–321.
- (50) Martinelli, L.; Jameson, A. Computational aerodynamics: solvers and shape optimization. *Journal of Heat Transfer* **2013**, *135*, 011002.
- (51) Seabra, M.; Azevedo, J.; Araújo, A.; Reis, L.; Pinto, E.; Alves, N.; Santos, R.; Mortágua, J. P.; et al. Selective laser melting (SLM) and topology optimization for lighter aerospace components. *Procedia Structural Integrity* **2016**, *1*, 289–296.
- (52) Wang, X.; Xu, S.; Zhou, S.; Xu, W.; Leary, M.; Choong, P.; Qian, M.; Brandt, M.; Xie, Y. M. Topological design and additive manufacturing of porous metals for bone scaffolds and orthopaedic implants: A review. *Biomaterials* **2016**, *83*, 127–141.
- (53) Al-Tamimi, A. A.; Almeida, H.; Bartolo, P. Structural optimization for medical implants through additive manufacturing. *Progress in Additive Manufacturing* **2020**, *5*, 95–110.
- (54) Lu, L.; Joannopoulos, J. D.; Soljačić, M. Topological photonics. *Nat. Photonics* **2014**, *8*, 821–829.
- (55) Piggott, A. Y.; Lu, J.; Lagoudakis, K. G.; Petykiewicz, J.; Babinec, T. M.; Vučković, J. Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer. *Nat. Photonics* **2015**, *9*, 374–377.
- (56) Lin, Z.; Liang, X.; Lončar, M.; Johnson, S. G.; Rodriguez, A. W. Cavity-enhanced second-harmonic generation via nonlinear-overlap optimization. *Optica* **2016**, *3*, 233–238.
- (57) Sell, D.; Yang, J.; Doshay, S.; Yang, R.; Fan, J. A. Large-angle, multifunctional metagratings based on freeform multimode geometries. *Nano Lett.* **2017**, *17*, 3752–3757.
- (58) Christiansen, R. E.; Lin, Z.; Roques-Carnes, C.; Salamin, Y.; Kooi, S. E.; Joannopoulos, J. D.; Soljačić, M.; Johnson, S. G. Fullwave Maxwell inverse design of axisymmetric, tunable, and multi-scale multi-wavelength metalenses. *Opt. Express* **2020**, *28*, 33854–33868.
- (59) Augenstein, Y.; Rockstuhl, C. Inverse design of nanophotonic devices with structural integrity. *ACS photonics* **2020**, *7*, 2190–2196.
- (60) Sigmund, O. On the usefulness of non-gradient approaches in topology optimization. *Structural and Multidisciplinary Optimization* **2011**, *43*, 589–596.
- (61) Hegde, R. S. Photonics inverse design: pairing deep neural networks with evolutionary algorithms. *IEEE J. Sel. Top. Quantum Electron.* **2020**, *26*, 1–8.
- (62) Ren, Y.; Zhang, L.; Wang, W.; Wang, X.; Lei, Y.; Xue, Y.; Sun, X.; Zhang, W. Genetic-algorithm-based deep neural networks for highly efficient photonic device design. *Photonics Research* **2021**, *9*, B247–B252.
- (63) Woldseth, R. V.; Aage, N.; Bærentzen, J. A.; Sigmund, O. On the use of artificial neural networks in topology optimization. *Structural and Multidisciplinary Optimization* **2022**, *65*, 294.
- (64) Peurifoy, J.; Shen, Y.; Jing, L.; Yang, Y.; Cano-Renteria, F.; DeLacy, B. G.; Joannopoulos, J. D.; Tegmark, M.; Soljačić, M. Nanophotonic particle simulation and inverse design using artificial neural networks. *Science Advances* **2018**, *4*, No. eaar4206.
- (65) Deng, Y.; Ren, S.; Fan, K.; Malof, J. M.; Padilla, W. J. Neural-adjoint method for the inverse design of all-dielectric metasurfaces. *Opt. Express* **2021**, *29*, 7526–7534.
- (66) Ren, S.; Padilla, W.; Malof, J. Benchmarking deep inverse models over time, and the neural-adjoint method. *arXiv:2009.12919 [cs.LG]* **2021**, na.
- (67) Augenstein, Y.; Repan, T.; Rockstuhl, C. Neural operator-based surrogate solver for free-form electromagnetic inverse design. *ACS Photonics* **2023**, *10*, 1547–1557.
- (68) Wu, C.; Jiao, Z.; Deng, H.; Huang, Y.-S.; Yu, H.; Takeuchi, I.; Rios Ocampo, C. A.; Li, M. Reconfigurable inverse-designed phase-change photonics. *APL Photonics* **2025**, *10*, 016113.
- (69) Xie, S.; Girshick, R.; Dollár, P.; Tu, Z.; He, K. Aggregated residual transformations for deep neural networks. *Proceedings of the IEEE conference on computer vision and pattern recognition*; IEEE, 2017; pp 1492–1500.
- (70) Wiecha, P. R.; Muskens, O. L. Deep Learning Meets Nanophotonics: A Generalized Accurate Predictor for Near Fields and Far Fields of Arbitrary 3D Nanostructures. *Nano Lett.* **2020**, *20*, 329–338.
- (71) Khaireh-Walieh, A.; Langevin, D.; Bennet, P.; Teytaud, O.; Moreau, A.; Wiecha, P. R. A Newcomer's Guide to Deep Learning for Inverse Design in Nano-Photonics. *Nanophotonics* **2023**, *12*, 4387–4414.
- (72) Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; Courville, A. C. Improved training of wasserstein gans. *Advances in Neural Information Processing Systems* **2017**, *30*, na.
- (73) Taballione, C.; van der Meer, R.; Snijders, H. J.; Hooijschuur, P.; Epping, J. P.; de Goede, M.; Kassenberg, B.; Venderbosch, P.; Toebes, C.; van den Vlekkert, H.; et al. A universal fully reconfigurable 12-mode quantum photonic processor. *Materials for Quantum Technology* **2021**, *1*, 035002.
- (74) Blanchard-Dionne, A.-P.; Martin, O. J. F. Successive Training of a Generative Adversarial Network for the Design of an Optical Cloak. *OSA Continuum* **2021**, *4*, 87–95.
- (75) de Goede, M.; Snijders, H.; Venderbosch, P.; Kassenberg, B.; Kannan, N.; Smith, D.; Taballione, C.; Epping, J. P.; van den Vlekkert, H.; Renema, J. J. High fidelity 12-mode quantum photonic processor operating at InGaAs quantum dot wavelength. *Integrated Photonics Research, Silicon and Nanophotonics* **2022**, ITu4B–3.
- (76) Taballione, C.; Anguita, M. C.; de Goede, M.; Venderbosch, P.; Kassenberg, B.; Snijders, H.; Kannan, N.; Vleeshouwers, W. L.; Smith, D.; Epping, J. P.; et al. 20-mode universal quantum photonic processor. *Quantum* **2023**, *7*, 1071.
- (77) Klokko, N.; Gorecki, J.; Beddoes, B.; Apostolopoulos, V. Deep neural network ensembles for THz-TDS refractive index extraction exhibiting resilience to experimental and analytical errors. *Opt. Express* **2023**, *31*, 44575–44587.
- (78) Nomura, K. Spin models constructed from Hadamard matrices. *Journal of Combinatorial Theory, Series A* **1994**, *68*, 251–261.
- (79) Leung, D. Simulation and reversal of n-qubit Hamiltonians using Hadamard matrices. *Journal of Modern Optics* **2002**, *49*, 1199–1217.
- (80) Tadej, W.; Życzkowski, K. A concise guide to complex Hadamard matrices. *Open Systems & Information Dynamics* **2006**, *13*, 133–177.
- (81) Frey, B. J.; Leviton, D. B.; Madison, T. J. Temperature-dependent refractive index of silicon and germanium. *Optomechanical Technologies for Astronomy* **2006**, 62732J.
- (82) Reck, M.; Zeilinger, A.; Bernstein, H. J.; Bertani, P. Experimental realization of any discrete unitary operator. *Physical review letters* **1994**, *73*, 58.

(83) Clements, W. R.; Humphreys, P. C.; Metcalf, B. J.; Kolthammer, W. S.; Walmsley, I. A. Optimal design for universal multiport interferometers. *Optica* **2016**, 3, 1460–1465.

(84) Blundell, S.; Radford, T.; Ajia, I. A.; Lawson, D.; Yan, X.; Banakar, M.; Thomson, D. J.; Zeimpekis, I.; Muskens, O. L. Ultracompact programmable silicon photonics using layers of low-loss phase-change material Sb_2Se_3 of increasing thickness. *arXiv:2409.12582 [physics.optics]* **2024**, na.