UC Irvine
UC Irvine Previously Published Works

Title
Arbitrary control over multimode wave propagation for machine learning

Permalink
https://escholarship.org/uc/item/10z3s93d

Journal
Nature Physics, 22(1)

ISSN
1745-2473

Authors
Onodera, Tatsuhiro
Stein, Martin M
Ash, Benjamin A
et al.

Publication Date
2026

DOI
10.1038/s41567-025-03094-2

Copyright Information
This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

nature physics

Article

https://doi.org/10.1038/s41567-025-03094-2

Arbitrary control over multimode wave propagation for machine learning

Received: 22 January 2025

Accepted: 7 October 2025

Published online: 8 December 2025

Check for updates

Tatsuhiro Onodera[1,2,8], Martin M. Stein[1,2,8] ✉, Benjamin A. Ash[1], Mandar M. Sohoni[1], Melissa Bosch[1], Ryotatsu Yanagimoto[1,2], Marc Jankowski[2,3], Timothy P. McKenna[2,3], Tianyu Wang[1,5], Gennady Shvets[1], Maxim R. Shcherbakov[1,6], Logan G. Wright[1,2,7] & Peter L. McMahon[1,4] ✉

Controlled multimode wave propagation can enable more space-efficient photonic processors than architectures based on discrete components connected by single-mode waveguides. Instead of defining discrete elements, one can sculpt the continuous substrate of a photonic processor to perform computations through multimode interference in two dimensions. Here we designed and demonstrated a device with a refractive index that can be rapidly reprogrammed across space, allowing arbitrary control of wave propagation. The device, a two-dimensional programmable waveguide, uses parallel electro-optic modulation of the refractive index of a slab waveguide with about 10^4 programmable spatial degrees of freedom. We implemented neural network inference on benchmark tasks with up to 49-dimensional vectors in a single pass, without digital pre-processing or post-processing. Theoretical and numerical analyses further indicated that two-dimensional programmable waveguides may offer not only a constant-factor reduction in device area but also a scaling benefit, with the area required growing as $N^{1.5}$ rather than N^2.

Deep neural networks (DNNs) have gained widespread adoption across many domains ranging from computer vision to natural language processing[1]. The size of DNN models has been increasing exponentially over the past decade, leading to exponentially increasing energy costs for running them. Limits to energy costs now impose a practical constraint on how large models can be[2], strongly motivating the exploration of alternative and energy-efficient computing approaches for executing DNNs, with a computational cost that is typically dominated by that of matrix–vector multiplications (MVMs). Optical neural networks, particularly those integrated on photonic chips that specialize in performing MVMs with optics instead of electronics, are one promising candidate approach[3–9].

The dominant paradigm for designing integrated photonic neural networks is to construct networks of discrete, programmable photonic components, such as Mach–Zehnder interferometers (MZIs), microring resonators or phase-change memory cells, connected by single-mode waveguides[10]. These networks execute a linear optical operation: an MVM between the vector encoded in the optical input to the chip and the matrix programmed in the discrete components. However, the maximum vector size, N, supported by chips using this approach has so far been restricted (Supplementary Table 1) to sizes far below what is necessary for optics to deliver an energy-efficiency advantage ($N \gtrsim 1{,}000$)[11–14]. The scale of such chips has been limited by at least two factors. First, the co-integration of optical and electrical components poses the challenge of routing $\mathcal{O}(N^2)$ electronic wires carrying the parameters of the matrix multiplication through the perimeter of the chip, often limiting the controllable degrees of freedom to a few hundred parameters. Second, individual optical components are rather

[1]School of Applied and Engineering Physics, Cornell University, Ithaca, NY, USA. [2]NTT Physics and Informatics Laboratories, NTT Research, Sunnyvale, CA, USA. [3]E. L. Ginzton Laboratory, Stanford University, Stanford, CA, USA. [4]Kavli Institute at Cornell for Nanoscale Science, Cornell University, Ithaca, NY, USA. [5]Present address: Department of Electrical and Computer Engineering, Boston University, Boston, MA, USA. [6]Present address: Department of Electrical Engineering and Computer Science, University of California, Irvine, CA, USA. [7]Present address: Department of Applied Physics, Yale University, New Haven, CT, USA. [8]These authors contributed equally: Tatsuhiro Onodera, Martin M. Stein. ✉e-mail: ms3452@cornell.edu; pmcmahon@cornell.edu

https://doi.org/10.1038/s41567-025-03094-2

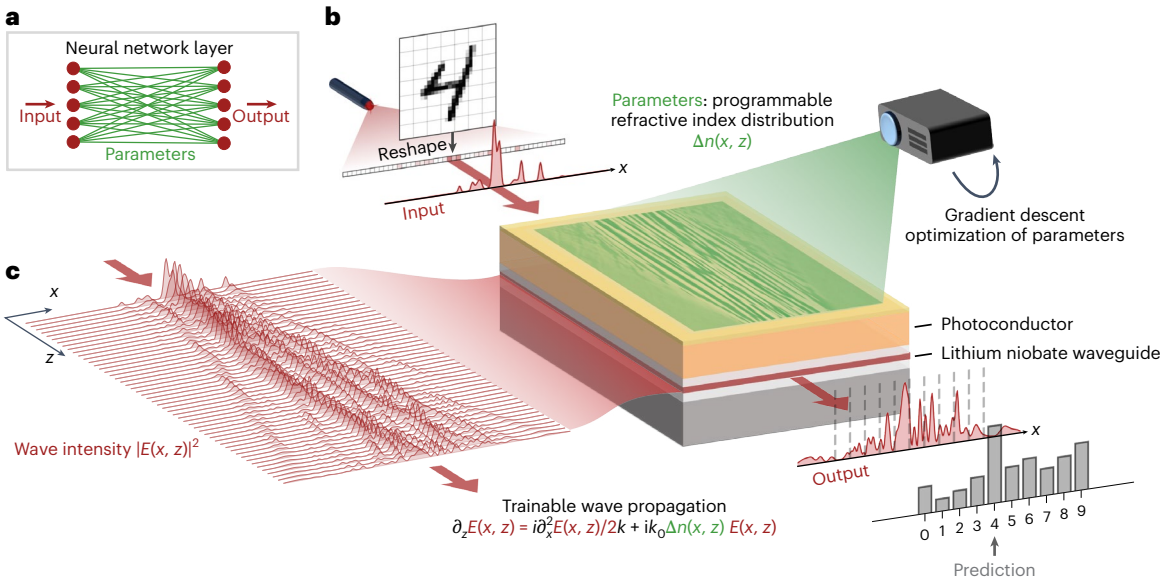


Fig. 1 | Machine learning with multimode wave propagation in the 2D-programmable waveguide. a, The fundamental unit of an artificial neural network (a layer) transforms an input vector into an output vector via a trainable matrix multiplication. **b**, Analogous to a neural network layer, the 2D-programmable waveguide linearly transforms an input optical field into an output optical field via wave propagation through a lithium niobate slab waveguide with 2D refractive index modulation $\Delta n(x, z)$ that can be continuously and arbitrarily programmed (up to practical limits on resolution and maximum modulation; Supplementary Information sections 1C,D). This refractive index modulation, which is directly set by the illumination pattern projected onto the device (shown in green), is trained so that wave propagation through the waveguide performs machine learning tasks (handwritten-digit classification shown here as an example). To determine the result of the classification, the intensity of the output beam is measured across equally sized bins; the bin with the highest total power corresponds to the predicted label. **c**, Simulated wave intensity in the slab waveguide, which shows that the neural network computation is performed with complex multimode wave propagation.

bulky owing to the comparably large optical wavelength and the often weak programmability of optical materials. On top of that, substantial portions of the chip's area need to be dedicated to sprawling interconnection regions comprising well-isolated waveguides.

We could achieve far greater spatial efficiency if we treated the entire chip as a blank slate with refractive index distribution, $n(x, z)$, as a function of space, which we could arbitrarily and reprogrammably sculpt, instead of building the integrated photonic neural network from discrete components[15–20]. This study addressed a central challenge: how to make a photonic chip with a programmable refractive index distribution, ideally without the integration complexity of introducing electronic wiring. In conventional nanophotonic chips, $n(x, z)$ is controlled by etching away material in lithographically defined regions and is fixed at fabrication time. Although inverse-designed chips[21] realizing MVMs with fixed matrices can be made[19], we generally would like to be able to program the matrix. Photorefractive crystals were explored several decades ago as a means to implement programmable linear operations with slab waveguides[22,23] but fell out of favour because the small achievable refractive index modulation (10^{-4}) meant that even centimetre-scale waveguides were unable to perform large-scale operations. Additionally, phase-change materials have recently been demonstrated to realize arbitrary refractive index distributions[24,25] but suffer from a limited number of rewrite cycles (4,000 in ref. 26) and high loss (greater than 2.8 dB mm^{-1} in refs. 25,26). The scale is currently limited; a previous study reported programming a device with a three-dimensional input and a three-dimensional output[25].

Here we introduce a photonic chip with a waveguide that is fully programmable in two dimensions: a two-dimensional (2D) programmable waveguide. The chip uses massively parallel electro-optic modulation to program $n(x, z)$ across ~10,000 individual regions of a lithium niobate slab waveguide. We trained multimode photonic structures within the chip that perform neural network inference (Fig. 1). The structures realized by our 2D-programmable waveguide are similar to inverse-designed nanophotonic devices[19,21]; they are computer-optimized 2D metastructures that control multimode wave propagation. A distinguishing feature of our device is its programmability, setting it apart from typical inverse-designed photonic devices, which are fixed after manufacturing. We achieve programmability optically, decoupling the electronic wiring for programming from the photonic chip; a pattern of light shone on top of our device creates a spatially varying refractive index modulation $\Delta n(x, z)$ in the slab waveguide. This is achieved by using the principle of photoconductive gain[27,28] to induce a refractive index modulation via the strong electro-optic effect in lithium niobate. In our study, we show how we can train the refractive index distribution so that the complex wave propagation through the device performs a desired neural network inference (Fig. 1c).

Operating principle of the device

Figure 1 shows a schematic of the machine learning process using the device. In machine learning, data are often encoded in vector form, $\mathbf{x}_l$, and processed via programmable matrices W and activation functions f in 'layers': $\mathbf{x}_{l+1} = f(W\mathbf{x}_l)$ (illustrated in Fig. 1a). We amplitude-encoded the machine learning input data $\mathbf{x}_l$ into the one-dimensional (1D) input optical field distribution, $E(x, z = 0)$ (Fig. 1b), which serves as the initial condition for programmable wave propagation described by the following partial differential equation:

$$\frac{\partial E(x,z)}{\partial z} = \frac{\mathrm{i}}{2k}\frac{\partial^2 E(x,z)}{\partial x^2} + \mathrm{i}k_0\Delta n(x,z)E(x,z), \qquad (1)$$

where z denotes the propagation direction, x represents the transverse direction, k_0 is the wavenumber in free space and k is the wavenumber in the slab waveguide. A solution of the equation is shown in Fig. 1c. The refractive index distribution of the slab waveguide, $n(x, z) = n_0 + \Delta n(x, z)$, has two contributions: a spatially uniform part, n_0, which is the refractive index of the waveguide when no programming light is

https://doi.org/10.1038/s41567-025-03094-2

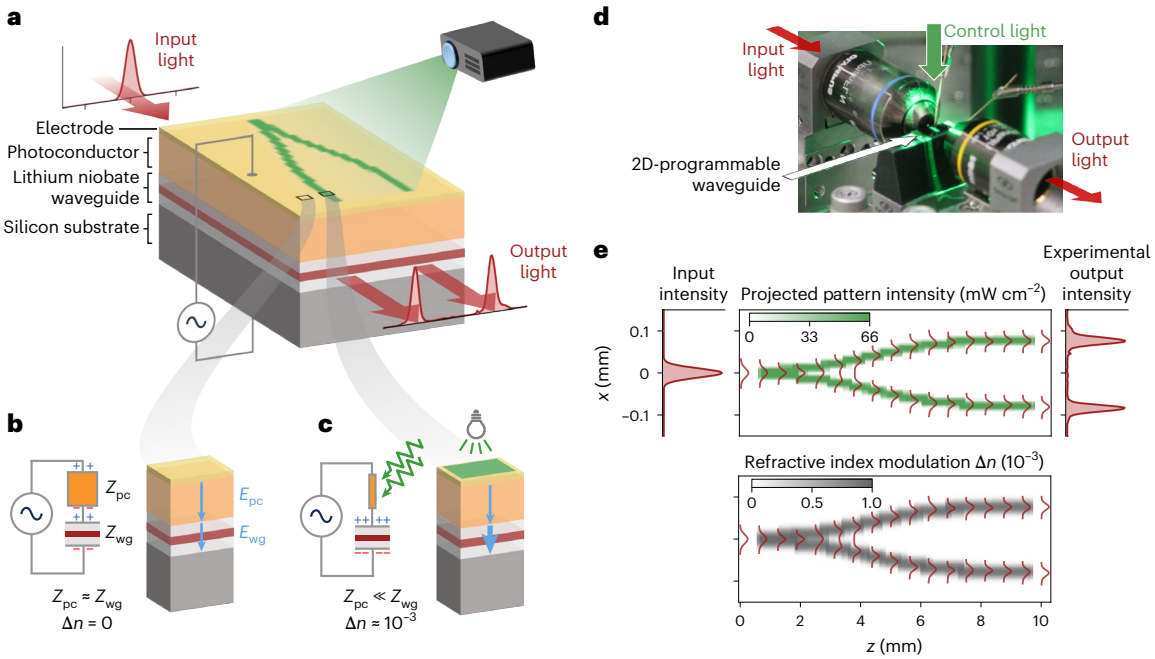


Fig. 2 | Operating principle of the 2D-programmable waveguide. a, The 2D-programmable waveguide consists of a nanophotonic stack of four layers: a conductive silicon substrate that doubles as the ground electrode, a Z-cut lithium niobate (in red) slab waveguide with silicon dioxide cladding (in white), a photoconductive layer for optical control of the refractive index and a gold electrode. **b**,**c**, Electrical circuit models of the 2D-programmable waveguide in regions with (**c**) and without (**b**) illumination. There is a voltage division between the photoconductor and the slab waveguide, with impedances, Z_{pc} and Z_{wg}, respectively. **c**, Upon illumination, the resistance of the photoconductor decreases, leading to an increase in the electric field (blue arrows) inside the waveguide, which induces a refractive index modulation via the electro-optic effect. **d**, A photograph of our prototype 2D-programmable waveguide in our experimental set-up. **e**, Experimental realization of a Y-branch splitter on the 2D-programmable waveguide, which splits the input light into two equal output beams. The projected pattern (in green) directly corresponds to the induced refractive index modulation (in grey). A simulation of the wave propagating through the device is overlaid with the patterns (in red).

impinging on it, and a programmable part that is induced by electro-optic modulation, $\Delta n(x, z)$, loosely corresponding to the programmable matrix weights W. After the optical field has propagated through the device, we measured its intensity, $I_{out}(x) \propto |E(x, z = L)|^2$, at the output facet and binned it to produce the output vector, $\mathbf{x}_{l+1}$, for the machine learning task.

Figure 2 shows how patterns of light program the refractive index modulation, $\Delta n(x, z)$. First, the device converts the projected pattern of light into a spatially varying quasi-static electric field. Second, the electro-optic waveguide changes its refractive index $\Delta n(x, z)$ in response to this electric field. Our device is an electro-optic waveguide coated by a photoconductive film and sandwiched between a pair of electrodes with an oscillating bias voltage applied across them. For regions of the chip that are illuminated at intensities of tens of milliwatts per square centimetre, the impedance of the photoconductor drops substantially, increasing the bias electric field within the slab waveguide. The lithium niobate waveguide locally changes its refractive index via the Pockels effect, owing to the spatially varying bias electric field. In our device, the largest refractive index modulation is approximately 10^{-3}, limited by the geometry of the material stack and a safety margin to prevent dielectric breakdown (Supplementary Information section 1C). The refractive index modulation in our device can take on continuous values by continuously varying the intensity of the projected pattern.

We projected illumination patterns across a 9 mm × 1 mm area, achieving a pixel resolution of 9 µm × 9 µm. (Limits to the spatial resolution are discussed in Supplementary Information section 1C.) This configuration enabled us to control the refractive index distribution, $n(x, z)$, with 10,000 degrees of freedom (see Supplementary Information section 3A for our calculation of the parameter count) and update the entire distribution at a rate of 3 Hz. To maximize the refractive index modulation, we set the amplitude of the oscillating bias voltage across the electrodes to be up to 1,000 V. Given that complementary metal–oxide–semiconductor electrode backplanes can only support spatially programmable voltages of around 10 V, our use of photoconductive gain was essential for achieving large electro-optic modulation. This approach allowed us to apply a large voltage to a single unpatterned electrode and realize controllable high voltages at virtual electrodes via the patterned illumination[27,28]. Owing to the large impedance of the device, less than 1 mW of electrical power was dissipated across its active area at the highest voltages.

To illustrate the operating principle of our device, we projected a pattern in the shape of a Y-branch splitter onto the 2D-programmable waveguide (Fig. 2a). The refractive index modulation was approximately proportional to the projected pattern, $\Delta n(x, z) \propto I(x, z)$ (Fig. 2e), up to spatial smoothing and a weak nonlinearity owing to voltage division. We coupled a single input Gaussian beam into the device using a beam shaper and measured the intensity of the output light with a camera, as shown in Fig. 2e. Because we used an oscillating bias voltage to drive the device, the induced refractive index pattern also oscillated in time. Therefore, measurements of the output intensity (and of the optical inputs, if not continuous wave) must coincide with the driving voltage peaking. We discuss approaches to overcome this limitation and further details on the experimental set-up in Methods.

Machine learning demonstrations with the 2D-programmable waveguide

We next applied the 2D-programmable waveguide to perform machine learning tasks: vowel classification[29] and Modified National Institute of Standards and Technology (MNIST) handwritten-digit

https://doi.org/10.1038/s41567-025-03094-2

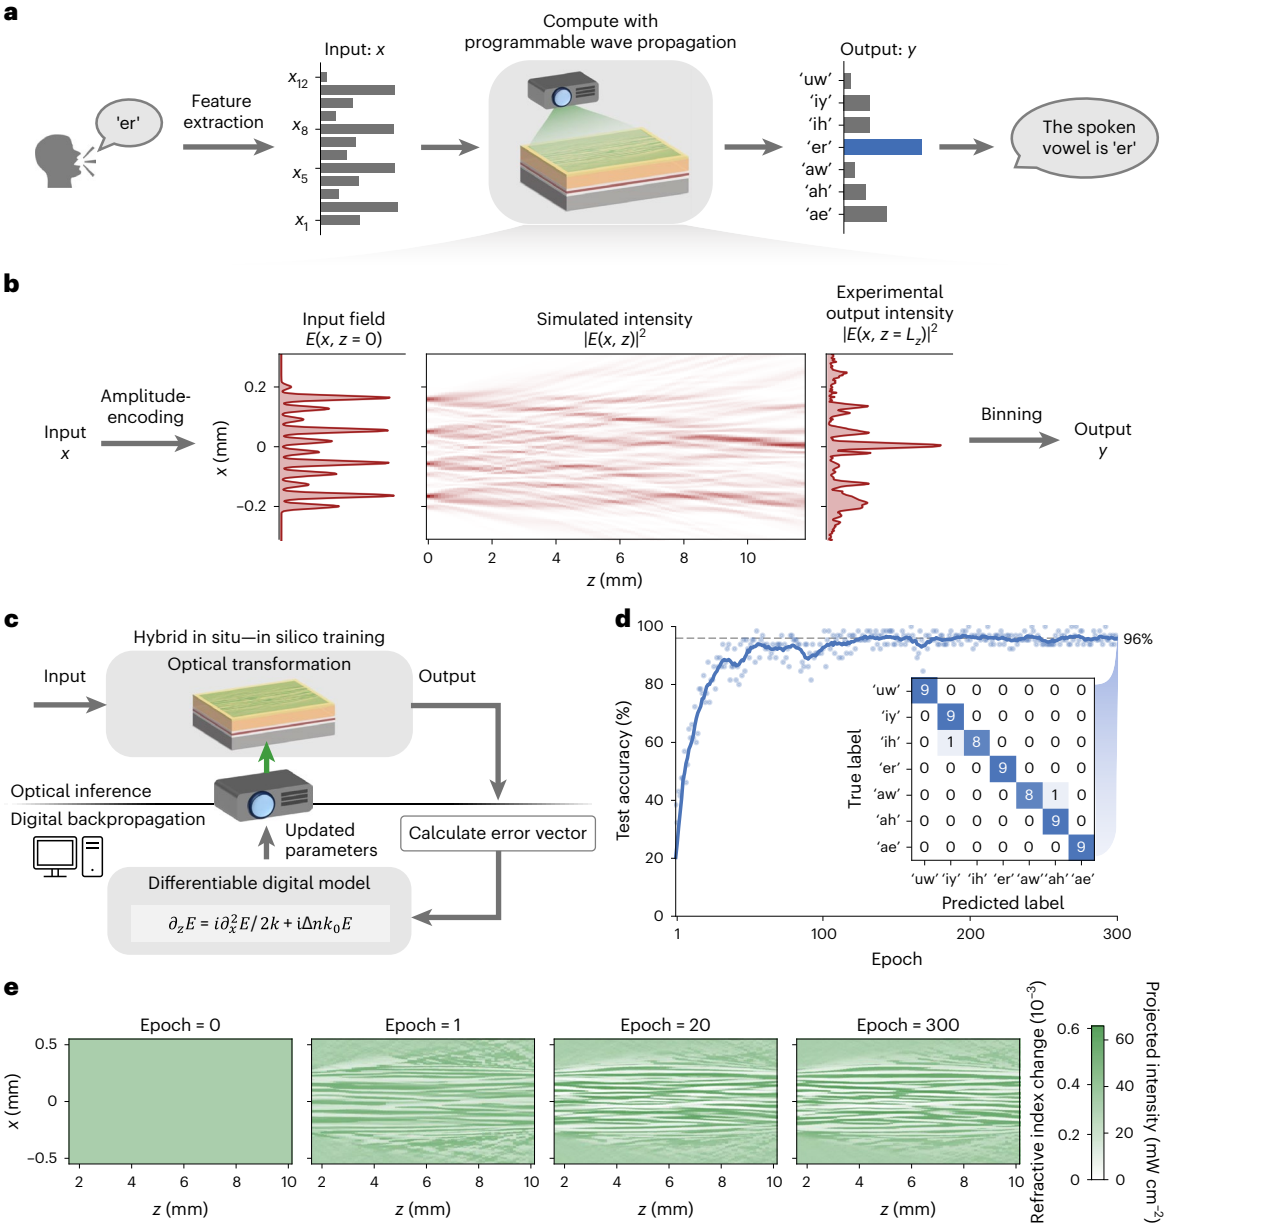


Fig. 3 | Vowel classification with the 2D-programmable waveguide. **a**, Overview of the approach. The task involved predicting a spoken vowel, here 'er', from a 12-dimensional input vector representing formant frequencies extracted from audio recordings. The 2D-programmable waveguide was trained to take in this input vector and output a seven-dimensional vector with a one-hot encoding format that indicates the predicted vowel. **b**, Left: the input vector was amplitude-encoded into 12 Gaussian spatial modes to produce the initial optical field distribution. Centre: the simulated wave propagation in the chip after training of the projected pattern. Right: the experimentally measured output intensity was binned by calculating the total power within equally sized spatial bins to produce the seven-dimensional output vector. **c**, Illustration of physics-aware training, a hybrid in situ–in silico backpropagation algorithm, which we used to train the parameters of the 2D-programmable waveguide. **d**, Test accuracy as a function of epoch. The inset shows the confusion matrix on the test dataset after training. **e**, Evolution of the trainable parameters (projected patterns) at different stages of training.

classification[30]. Both tasks are used as benchmarks in studies of similar on-chip optical neural networks, providing useful points of comparison[3,6,20].

The vowel classification dataset[29] comprises formant frequencies extracted from audio recordings of spoken vowels by various speakers. The task was to predict which of the seven vowels is spoken, given a 12-dimensional input vector of formant frequencies. We divided the dataset into a training set and a test set, comprising 196 (75% of the dataset) and 63 (25%) samples, respectively.

Figure 3 presents our experimental results on performing vowel classification with the 2D-programmable waveguide. As shown in Fig. 3a,b, we encoded the 12-dimensional input vectors in the amplitudes of 12 equidistant spatial Gaussian modes that were simultaneously projected onto the input facet of the device. For readout, we measured the intensity at the output facet with a camera and binned the camera pixels into seven different regions, with each region corresponding to a specific vowel. The predicted vowel for an input was given by the region that received the most optical power. Thus, as shown in the simulated intensity distribution, $|E(x,z)|^2$, in Fig. 3b, the device learned to use complex multimode wave propagation to direct most power towards the region corresponding to the correct vowel. Further details on output decoding and the overall computational model of our optical neural network demonstrations are provided in Supplementary Information section 5A.

https://doi.org/10.1038/s41567-025-03094-2

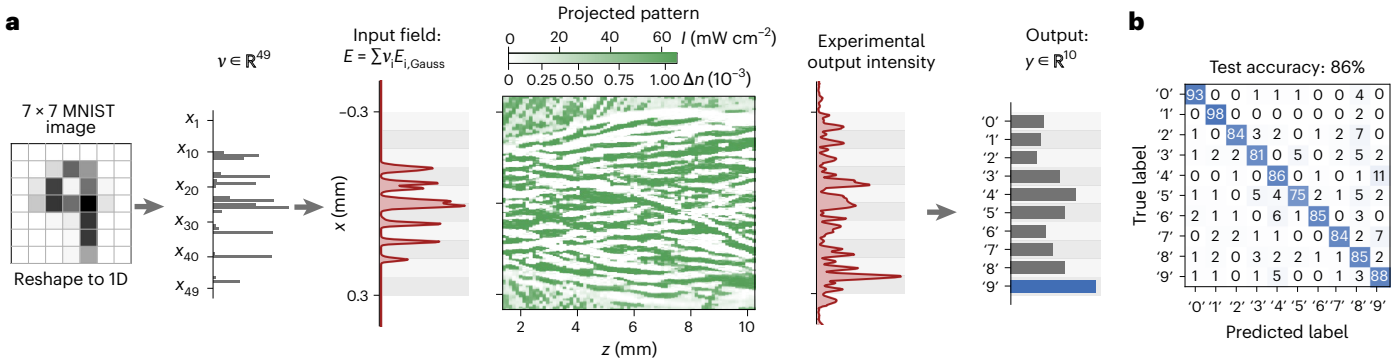


Fig. 4 | MNIST handwritten-digit classification with the 2D-programmable waveguide: neural network inference with high-dimensional input vectors. a, We performed MNIST handwritten-digit classification with the 2D-programmable waveguide. Each image from the MNIST dataset was electronically downsampled and reshaped to a 49-dimensional vector. We trained the device to perform machine learning on this high-dimensional input vector with the same procedure as the vowel classification task (Fig. 3). **b**, The confusion matrix was derived from evaluations on the test dataset comprising 10,000 images. After ten epochs of training, the system achieved 86% accuracy on the test dataset. As a baseline, a single-layer digital neural network with a 49 × 10 matrix achieves 90% accuracy on this same task.

The refractive index distribution to implement vowel classification was learned using physics-aware training[31], a modified backpropagation algorithm (Fig. 3c). In this algorithm, the forward pass is performed by the experimental set-up, whereas the backward pass is computed with a digital model of the experiment. The hybrid in situ–in silico nature of the algorithm allows for efficient training even in the presence of both imperfect models and experimental noise (Supplementary Information section 3C). The digital model was challenging to construct because of the large number of parameters and the complexity of the wave propagation. Initially, a purely physics-based model (using equation (1)) provided qualitative but not quantitative agreement with the experimental results. The remaining discrepancies were largely removed with data-driven refinements to the physics-based model (Supplementary Information section 4).

Using physics-aware training, we trained the 2D-programmable waveguide for a total of 300 epochs, which took approximately 1 h on the experimental set-up (Fig. 3d). Figure 3e shows the evolution of the initially uniform illumination pattern into a complex pattern that resembles the refractive index distributions found in inverse-designed photonic devices. Figure 3d shows that despite the complexity of the projected pattern, the 2D-programmable waveguide successfully performed the vowel classification task and achieved a test accuracy of 96% after training.

Figure 4 presents the experimental results on MNIST handwritten-digit classification. The task consisted of classifying 14-by-14-pixel images of handwritten digits from 0 to 9. We divided the MNIST dataset in the standard manner into 60,000 training images and 10,000 test images. We downsampled each MNIST image to 7-by-7 pixels and then flattened it to a 49-dimensional input vector.

To train the refractive index distribution to perform MNIST classification, we followed the same procedure used for the vowel classification task (Fig. 3). The 2D-programmable waveguide processed the 49-dimensional input vector to produce a ten-dimensional output vector that corresponds to the ten possible digits (more details in Supplementary Information section 5C). As shown in Fig. 4b, the system achieved 86% accuracy on the test dataset after ten epochs of training, which took about 10 h on the experimental set-up. This falls 4% short of the 90% accuracy that a one-layer digital neural network achieves on this downsampled MNIST classification task, likely because of imperfect modelling and experimental drifts. Nevertheless, this result demonstrates that complex wave propagation in our device can be harnessed to perform computations comparable to that of a single-layer neural network with a 49 × 10 matrix of trainable parameters.

Length (and area) scaling for 2D-programmable waveguides

We turn to a discussion of the theoretical size scaling of different on-chip photonic processors. Integrated photonics can be used to implement matrices of size $N \times N$ using circuits of width much wider than ~$N\lambda_0$ (N well-isolated waveguides, where λ_0 is the free-space wavelength) and length greater than ~$N\lambda_0/\Delta n_{\mathrm{prog}}$ (N π-phase shifters)[32,33]. Because the programmable refractive index Δn_{prog} is often very small, this limits circuits to be either very long or operate on low-dimensional inputs. In the 2D-programmable-waveguide paradigm presented in this paper, the required device width scales as ~$N\lambda_0$, imposed by the diffraction limit. One might intuitively expect that the length L_z of 2D-programmable waveguides would also need to scale as ~$N\lambda_0/\Delta n_{\mathrm{prog}}$, just as it does for the aforementioned circuit approaches for universal linear transformations on chip[32]. In this section, we present the analytical and numerical results showing that, beyond the constant-factor improvement in spatial footprint by avoiding single-mode waveguides, 2D-programmable waveguides may offer a different size scaling.

Our analytical argument, presented in detail in Supplementary Information section 8A,B, investigates the amount of phase shift necessary to optically implement a given unitary transformation of dimension N in a programmable multimode waveguide. As the unitary dimension increases, the phase shift per element, linked to the generator of the unitary, becomes smaller. (More precisely, for any given unitary, one can always find a generator with elements that have a root-mean-square (r.m.s.) smaller than $\pi/\sqrt{N}$). This suggests an exciting possibility. High-dimensional optical matrix–vector multipliers based on interference may require much less propagation distance through phase shifters than commonly assumed. This property of unitary matrices has recently[34] been exploited to show that in three-MZI meshes, the average length of phase shifters can scale as $1/\sqrt{N}$, but universal MZI meshes by construction have a circuit depth that scales as N, independent of how short the individual phase-shifter elements are. However, in 2D-programmable waveguides, this insight can potentially lead to large practical benefits because the devices effectively consist entirely of phase shifters, with nothing else constraining the total system length. The intuitive core of the mathematical argument is that if one needed a full π-phase shift per phase shifter, the required device length would scale as $N\lambda_0/\Delta n_{\mathrm{prog}}$, but because the required phase shift goes as $\pi/\sqrt{N}$, the total required length scales as $N\lambda_0/(\sqrt{N}\Delta n_{\mathrm{prog}}) = \sqrt{N}\lambda_0/\Delta n_{\mathrm{prog}}$.

We analytically found a refractive index distribution, $\Delta n_{\mathrm{prog}}(x, z)$, with a magnitude (or, equivalently, length) that scales only as $\sqrt{N}$ (as measured by the r.m.s. of the distribution). Our construction is similar

https://doi.org/10.1038/s41567-025-03094-2

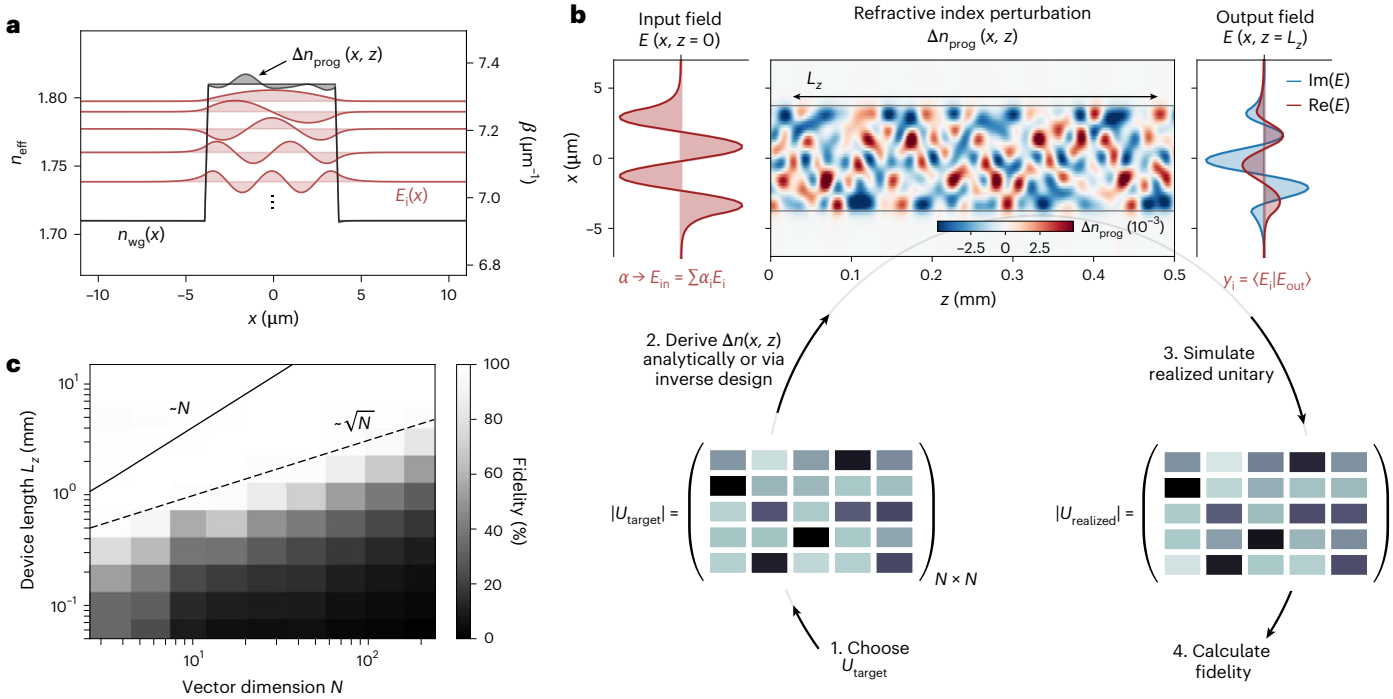


Fig. 5 | Simulation study showing the required length for devices realizing arbitrary photonic $N \times N$ unitaries scales as $\sqrt{N}$. a, We considered a weak refractive index profile $\Delta n_{prog}(x, z)$ embedded into a step-index multimode waveguide with profile $n_{wg}(x)$. As light propagated through the perturbed waveguide, its N eigenmodes interacted with each other, as described by the coupled-mode theory. These interactions generated an $N \times N$ unitary MVM. **b**, Vectors **a** are encoded in the modal amplitudes of the input electric field. Given a unitary matrix U_{target}, we analytically determined the refractive index perturbation $\Delta n_{prog}(x, z)$ embedded in the step-index waveguide (Supplementary Information section 8B) or via in silico inverse design (Supplementary Information section 8C). We verified that after propagation through the multimode waveguide, the modal amplitudes transformed according to $U_{realized} \approx U_{target}$. **c**, Fidelities of $N \times N$ Haar-random unitary matrices realized in waveguides of various lengths L_z with an inverse-designed $\Delta n_{prog}(x, z) < \Delta n_{max} = 5 \cdot 10^{-3}$. The separation between unitaries realized with high versus low fidelity is well described by the line $L_z = \sqrt{N}\lambda_0/\Delta n_{max}$.

to the one presented by Larocque and Englund[17] in that it couples the modes of a multimode waveguide but crucially differs by implementing parallel global couplings rather than a sequence of pairwise mode couplings, thereby achieving better scaling.

Our analytical argument relies on strong approximations. We used coupled-mode theory under a rotating wave approximation to calculate the propagation of unidirectional scalar waves in a perturbed multimode waveguide. To validate our theory, we present the results from numerical simulations that make far fewer assumptions. We simulated unidirectional scalar waves in a perturbed multimode waveguide (Supplementary Information section 8D). Our analytical argument only shows that the r.m.s. of the refractive index distribution scales as $\sim \sqrt{N}$. However, in practical devices, it is usually the maximum value of the refractive index change that is limited. Our simulations show that even with a strictly imposed maximum value of the programmable refractive index strength, unitaries can be implemented accurately over a propagation distance that scales as $\sim \sqrt{N}$. We emphasize that our analytical argument permits completely general dense unitary matrices.

This length scaling is surprising in part because fundamental geometric considerations in linear optical devices suggest that the length required to perform arbitrary operations on N modes is proportional to $N\lambda_0$ (refs. 35,36). Our theory, showing that $L_z \sim \sqrt{N}\lambda_0/\Delta n_{prog}$ is sufficient to perform universal optical operations on N modes, implies a different length scaling but does not contradict the bound by Miller[35] unless N is very large, often ≫1,000 using common values (Supplementary Information section 8D). Therefore, our result suggests that there is a large practical regime in which the better-than-linear length scaling can be exploited to create universal optical processors with modest programmability or short propagation distances.

However, we note that the length-scaling argument does not directly apply to the device used in our experiments because our scaling result assumes guided modes in the x direction, whereas our experiments were performed in a slab waveguide wide enough that light was effectively unguided. It is nonetheless the case that one could, as far as we can tell without having explicitly performed these experiments, realize devices and experiments that satisfy the assumptions of the theory (see discussion in Supplementary Information section 8D). For example, we simulated a multimode waveguide with a modest width of 0.3 mm, length of 5 mm and a refractive index programmability no larger than $\Delta n_{prog} = 5 \times 10^{-3}$ with 500-nm resolution, and showed that such a waveguide can realize a 200 × 200-dimensional unitary with high fidelity (Fig. 5 and Supplementary Information section 8D). A device measuring 1.5 mm in width and 11 mm in length, with the same programmable refractive index magnitude and resolution, should be able to realize arbitrary unitaries with dimensions as high as 1,000 × 1,000.

Discussion and outlook

We introduced and demonstrated a 2D-programmable photonic processor comprising a lithium niobate slab waveguide with a refractive index distribution, $n(x, z)$, that can be continuously programmed. The device design enables programming by parallel electro-optic modulation with approximately 10,000 degrees of freedom. We used our chip to perform neural network inference by training the refractive index distribution and consequently the multimode wave propagation through the chip. To train the device, we developed a physics-based model of the behaviour of the chip, along with a data-driven refinement, allowing the model to be sufficiently accurate that it supports backpropagation-based training[31].

The predominant approach to building integrated photonic neural networks is to fabricate large arrays of discrete components

connected by single-mode waveguides[10]. In contrast, we adopted the conceptual approach of using wave propagation in distributed spatial modes[15–19,23,37] and experimentally validated the theoretical predictions[16,17,37] that this approach will be more space-efficient. Our prototype chip was able to perform neural network inference with input vectors of dimension up to 49, which is larger than the capability of the neural network photonic chips reported in previous studies[3,6,8,9,20,38–40] and more space efficient than any of these chips based on networks of discrete components (see Supplementary Information section 6 for a detailed comparison). This large input dimension enabled us to use our chip to perform MNIST handwritten-digit classification with a single pass through the chip and without using any digital electronic parameters.

One of the most promising applications of our approach is in reducing the energy cost of neural network inference, which remains (and is likely to remain) dominated by linear MVM. There exists a breakeven point beyond which optical devices could outperform electronic hardware on this metric, owing to the more favourable energy scaling with dimension N (optics: $E \sim N$; electronics: $E \sim N^2$) (refs. 11–14). Reaching this regime requires high-dimensional operations. At low N, the overhead cost from analogue-to-digital conversion and optical–electronic transduction outweighs the benefits, with estimates placing the breakeven point around $N = 1{,}000$, far beyond what is currently possible on a single chip. We derived a theoretical scaling law (Supplementary Information section 8) describing how the dimension of possible MVMs in a 2D-programmable waveguide scales with the device dimensions and the refractive index change. Surprisingly, the device length only needs to scale as $\sqrt{N}$, better than the most common approaches of designing photonic circuits[32,33]. This result may enable all-optical matrix–vector multipliers with a dimension exceeding the breakeven point of energy efficiency. The development of such devices would make hybrid neural network architectures, in which analogue optics performs the linear operations and electronic circuits implement the nonlinearities, energy-competitive, changing the energy scaling of neural network inference.

To conclude, we believe that our device concept, with its ability to programmably control multimode wave propagation, may create new opportunities in the broader fields of optical computing and optical information processing[10,14,41]. Although our study has focused on machine learning, our device could also be used to solve integral equations[42] and combinatorial optimization problems[43]. More broadly, our chip is essentially an arbitrary (passive) photonic device that can be reconfigured on demand. Any photonic device that can be specified as an inhomogeneous refractive index distribution can be realized. Such devices can even be learned directly and effectively by performing inverse design[21] but in situ in real time. Our concept will potentially enable the development of reprogrammable photonic simulators supporting new studies of bound states in the continuum[44] and topological photonics[45] and applications in engineering. It may ultimately even be possible to make a device that combines programmable linear wave propagation (this study), programmable nonlinear wave propagation[15,18,46,47] (a natural extension of this study to having programmable $\chi^{(2)}(x, z)$ (ref. 48)) and programmable gain/loss (demonstrated in ref. 20), giving rise to a reconfigurable on-chip platform capable of realizing almost every functionality we have in free-space optics.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41567-025-03094-2.

References

1. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
2. Patterson, D. et al. Carbon emissions and large neural network training. Preprint at https://arxiv.org/abs/2104.10350 (2021).
3. Shen, Y. et al. Deep learning with coherent nanophotonic circuits. *Nat. Photon.* **11**, 441–446 (2017).
4. Tait, A. N. et al. Neuromorphic photonic networks using silicon photonic weight banks. *Sci. Rep.* **7**, 7430 (2017).
5. Feldmann, J., Youngblood, N., Wright, C. D., Bhaskaran, H. & Pernice, W. H. All-optical spiking neurosynaptic networks with self-learning capabilities. *Nature* **569**, 208–214 (2019).
6. Feldmann, J. et al. Parallel convolutional processing using an integrated photonic tensor core. *Nature* **589**, 52–58 (2021).
7. Huang, C. et al. A silicon photonic–electronic neural network for fibre nonlinearity compensation. *Nat. Electron.* **4**, 837–844 (2021).
8. Ashtiani, F., Geers, A. J. & Aflatouni, F. An on-chip photonic deep neural network for image classification. *Nature* **606**, 501-506 (2022).
9. Bandyopadhyay, S. et al. Single-chip photonic deep neural network with forward-only training. *Nat. Photon.* **18**, 1335–1343 (2024).
10. Shastri, B. J. et al. Photonics for artificial intelligence and neuromorphic computing. *Nat. Photon.* **15**, 102–114 (2021).
11. Hamerly, R., Bernstein, L., Sludds, A., Soljačić, M. & Englund, D. Large-scale optical neural networks based on photoelectric multiplication. *Phys. Rev. X* **9**, 021032 (2019).
12. Nahmias, M. A. et al. Photonic multiply-accumulate operations for neural networks. *IEEE J. Sel. Top. Quantum Electron.* **26**, 1–18 (2019).
13. Anderson, M. G., Ma, S.-Y., Wang, T., Wright, L. G. & McMahon, P. L. Optical transformers. Preprint at https://arxiv.org/abs/2302.10360 (2023).
14. McMahon, P. L. The physics of optical computing. *Nat. Rev. Phys.* **5**, 717–734 (2023).
15. Hughes, T. W., Williamson, I. A. D., Minkov, M. & Fan, S. Wave physics as an analog recurrent neural network. *Sci. Adv.* **5**, eaay6946 (2019).
16. Khoram, E. et al. Nanophotonic media for artificial neural inference. *Photonics Res.* **7**, 823–827 (2019).
17. Larocque, H. & Englund, D. Universal linear optics by programmable multimode interference. *Opt. Express* **29**, 38257–38267 (2021).
18. Nakajima, M., Tanaka, K. & Hashimoto, T. Neural Schrödinger equation: physical law as deep neural network. *IEEE Trans. Neural Networks Learn. Syst.* **33**, 2686–2700 (2022).
19. Nikkhah, V. et al. Inverse-designed low-index-contrast structures on a silicon photonics platform for vector–matrix multiplication. *Nat. Photon.* **18**, 501–508 (2024).
20. Wu, T., Menarini, M., Gao, Z. & Feng, L. Lithography-free reconfigurable integrated photonic processor. *Nat. Photon.* **17**, 710–716 (2023).
21. Molesky, S. et al. Inverse design in nanophotonics. *Nat. Photon.* **12**, 659–670 (2018).
22. Psaltis, D., Brady, D., Gu, X.-G. & Lin, S. Holography in artificial neural networks. *Nature* **343**, 325–330 (1990).
23. Brady, D. J. & Psaltis, D. Holographic interconnections in photorefractive waveguides. *Appl. Opt.* **30**, 2324–2333 (1991).
24. Delaney, M. et al. Nonvolatile programmable silicon photonics using an ultralow-loss Sb_2Se_3 phase change material. *Sci. Adv.* **7**, eabg3500 (2021).
25. Wu, C. et al. Freeform direct-write and rewritable photonic integrated circuits in phase-change thin films. *Sci. Adv.* **10**, eadk1361 (2024).

26. 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: Sb_2S_3 and Sb_2Se_3. *Adv. Funct. Mater.* **30**, 2002447 (2020).
27. Chiou, P. Y., Ohta, A. T. & Wu, M. C. Massively parallel manipulation of single cells and microparticles using optical images. *Nature* **436**, 370–372 (2005).
28. Wu, M. C. Optoelectronic tweezers. *Nat. Photon.* **5**, 322–324 (2011).
29. Hillenbrand, J., Getty, L. A., Clark, M. J. & Wheeler, K. Acoustic characteristics of American english vowels. *J. Acoust. Soc. Am.* **97**, 3099–3111 (1995).
30. LeCun, Y. *The MNIST Database of Handwritten Digits* (New York Univ., accessed 07 January 2025); http://yann.lecun.com/exdb/mnist/
31. Wright, L. G. et al. Deep physical neural networks trained with backpropagation. *Nature* **601**, 549–555 (2022).
32. Wu, B., Zhou, H., Dong, J. & Zhang, X. Programmable integrated photonic coherent matrix: principle, configuring, and applications. *Appl. Phys. Rev.* **11**, 011309 (2024).
33. Clements, W. R., Humphreys, P. C., Metcalf, B. J., Kolthammer, W. S. & Walmsley, I. A. Optimal design for universal multiport interferometers. *Optica* **3**, 1460–1465 (2016).
34. Hamerly, R., Basani, J. R., Sludds, A., Vadlamani, S. K, and Englund, D. To- ward the information-theoretic limit of programmable photonics. *APL Photonics* **10**, 110803 (2025).
35. Miller, D. A. B. Why optics needs thickness. *Science* **379**, 41–45 (2023).
36. Li, S. & Hsu, C. W. Thickness bound for nonlocal wide-field-of-view metalenses. *Light Sci. Appl.* **11**, 338 (2022).
37. Gu, J. et al. M3icro: Ma- chine learning-enabled compact photonic tensor core based on programmable multi-operand multimode interference. *APL Mach. Learn.* **2**, 016106 (2024).
38. Huang, C. et al. Demonstration of scalable microring weight bank control for large-scale photonic integrated circuits. *APL Photonics* **5**, 040803 (2020).
39. Zhang, H. et al. An optical neural chip for implementing complex-valued neural network. *Nat. Commun.* **12**, 457 (2021).
40. Pai, S. et al. Experimentally realized in situ backpropagation for deep learning in photonic neural networks. *Science* **380**, 398–404 (2023).
41. Wetzstein, G. et al. Inference in artificial intelligence with deep optics and photonics. *Nature* **588**, 39–47 (2020).
42. Mohammadi Estakhri, N., Edwards, B. & Engheta, N. Inverse-designed metastructures that solve equations. *Science* **363**, 1333–1338 (2019).
43. Roques-Carmes, C. et al. Heuristic recurrent algorithms for photonic ising machines. *Nat. Commun.* **11**, 249 (2020).
44. Hsu, C. W., Zhen, B., Stone, A. D., Joannopoulos, J. D. & Soljačić, M. Bound states in the continuum. *Nat. Rev. Mater.* **1**, 1–13 (2016).
45. Price, H. et al. Roadmap on topological photonics. *J. Phys. Photonics* **4**, 032501 (2022).
46. Marcucci, G., Pierangeli, D. & Conti, C. Theory of neuromorphic computing by waves: machine learning by rogue waves, dispersive shocks, and solitons. *Phys. Rev. Lett.* **125**, 093901 (2020).
47. Teğin, U., Yıldırım, M., Oğuz, İ., Moser, C. & Psaltis, D. Scalable optical learning operator. *Nat. Comput. Sci.* **1**, 542–549 (2021).
48. Yanagimoto, R. et al. Programmable on-chip nonlinear photonics. *Nature* (2025).

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

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

© The Author(s) 2025

Methods

Device fabrication and design

As shown in Extended Data Fig. 1, we started our fabrication processes from a thin-film lithium niobate wafer purchased from NANOLN. It was a p-type silicon wafer with a substrate conductivity of 0.01–0.05 Ω cm, 2 µm of silicon dioxide deposited via plasma-enhanced chemical vapour deposition (PECVD) and 700 nm of Z-cut MgO-doped lithium niobate that was wafer-bonded with the ion-cut technique. We diced small pieces from the wafer using a DISCO Dicing Saw for further processing. We deposited an additional 1 µm of silicon dioxide via PECVD as a cladding, followed by another deposition of 4 µm of silicon-rich silicon nitride, which was the photoconductive layer, via PECVD. The silicon-rich silicon nitride layer was deposited in an Oxford Plasmalab 100 by flowing 40 sccm of SiH_4, 10 sccm of N_2O and 1,425 sccm of N_2 into the deposition chamber at a temperature of 350 °C and a pressure of 1,900 mTorr. We alternated pulses of high-frequency and low-frequency power during deposition to minimize film stress using 160-W low-frequency pulses for 12 s and 200-W high-frequency pulses for 8 s.

Next, we evaporated electrodes onto the chip using a CVC SC4500 E-gun Evaporation System. We first evaporated 10 nm of titanium as an adhesion layer and then 5 nm of gold. To prevent dielectric breakdown between the top electrode and the conductive substrate through air at the edges of the chip, we covered the perimeter of the chip with tape before evaporation. The tape acted as a mask, preventing deposition closer than around 1 mm to the edges, thereby increasing the path length between the top electrode and substrate through air. To minimize coupling losses into the waveguide, we used an Allied MultiPrep Polisher to polish the waveguide facets. We polished using silicon carbide paper of successively finer grain size, starting at 3 µm and then moving to 1 µm and 0.5 µm.

Device design and characterization

We used Z-cut lithium niobate for the slab waveguide so that the crystal axis of the lithium niobate is parallel to the strong electric field induced by the bias voltage, which points out of the waveguide plane (in the y direction). As shown in Supplementary Fig. 1, we used the transverse magnetic mode of the slab waveguide, with an optical electric field that is also oriented in the y direction. Because the r_{33} electro-optic coefficient is largest in lithium niobate, this configuration maximized the strength of the electro-optic modulation.

The thickness of the lithium niobate layer was chosen for single-mode operation (Supplementary Information section 1A). To maximize the refractive index modulation, it is beneficial to have a thicker photoconductor and a thinner silicon dioxide cladding (Supplementary Information section 1B). The silicon dioxide cladding was chosen to be sufficiently thick to ensure low propagation loss. Thus, we balanced these tradeoffs to arrive at the device geometry shown in Extended Data Fig. 1. The device has a propagation loss of <1 dB cm^{-1} at a wavelength of 1,550 nm (Supplementary Information section 1E).

To maximize the refractive index contrast of the slab waveguide between the bright (illuminated) and dark regions, it is important to design the photoconductor to have high dark resistance and low bright resistance. Thus, to optimize the material properties of the photoconductor, we swept the silicon-to-nitrogen ratio in the silicon-rich silicon nitride photoconductor (by varying the amount of SiH_4 gas we flowed into the PECVD) and chose the material with the largest photoconductive contrast. We characterized the refractive index modulation as a function of the intensity of the projected pattern with an off-axis holography set-up (Supplementary Information section 1C). The maximum refractive index modulation that we achieved in this study was approximately 10^{-3}. We show in Supplementary Fig. 3 that this can be increased to beyond 4×10^{-3} by using a photoconductor layer that is twice as thick (8 µm) and by further optimizing the photoconductive properties. In Supplementary Information section 7B, we also discuss how the refractive index modulation can be further increased by switching to a different material for the waveguide core.

Experimental set-up

As shown in Extended Data Fig. 2, the experimental set-up can be roughly divided into five units: (1) an optical beam shaper to create spatially varying 1D electric field inputs for the 2D-programmable waveguide; (2) a projector to create a programmable illumination pattern that controls the refractive index distribution inside the waveguide; (3) a butt-coupling set-up to couple light in and out of the 2D-programmable waveguide; (4) a high-voltage source to apply an oscillating bias voltage across the electrodes of the 2D-programmable waveguide; and (5) a camera to measure the intensity of the output beam. We note that the experimental set-up relies on more free-space optical components than usual for an integrated photonics experiment. This is a direct consequence of our decision to keep the fabrication of the device simple, without lithographically defined structures for this initial proof-of-concept demonstration. We envision that a more compact, fully integrated version of the 2D-programmable waveguide could be built by integrating on-chip lithium niobate modulators and detectors and a micro-light-emitting-diode display (Supplementary Information section 7D).

In this section, we provide an overview of the key components and functionalities of the experimental set-up. For a more detailed description, including photographs of the optical set-up and specifics on the components, such as part numbers and manufacturers, see Supplementary Information section 2.

The free-space beam shaper allows for the realization of arbitrary input optical fields $E(x, z = 0)$ up to a spatial resolution of 2 µm and over a distance of 600 µm. In this experiment, we used the beam shaper to create both simple input fields, such as a single Gaussian beam for the Y-branch splitter demonstration, and more complex input fields for the machine learning demonstrations. This flexibility also enabled us to freely vary the encoding of input vectors into the optical field. For instance, we varied the width of the input modes and adjusted their spacing, which was tailored to each machine learning task. Finally, because the beam shaper is capable of shaping both the amplitude and phase of the input field, it was also used to calibrate the 2D-programmable waveguide (Supplementary Information section 4B).

The design of the beam shaper that we built closely follows that of a previous study[49], which also programmably shapes the input light coupled into slab waveguides. The core working principle of the beam shaper is to create spatially varying phase gratings on a 2D-phase spatial light modulator[50] (SLM; Meadowlark Optics UHSP1K-850-1650-PC8). We varied the amplitude and relative positions of these phase gratings to control the input optical field $E(x, z = 0)$. A lens after the SLM performed a Fourier transform that separated the diffraction maxima of the phase gratings, and a spatial filter selected the first-order diffraction maximum. Finally, an optical relay system (comprising lenses 2 and 3, as shown in Extended Data Fig. 2) demagnified the optical field at the focal plane of the first lens and coupled the light into the 2D-programmable waveguide. Owing to the response time of the liquid crystal in SLM, the fastest update speed of the beam shaper was approximately 50 Hz.

The projector set-up was designed to create a high-resolution programmable illumination pattern over a large field of view. We used a digital micromirror device (DMD; Vialux V-7000) with a resolution of 1,024 × 768 pixels and a pixel pitch of 13.7 µm. The DMD was illuminated with green light (525 nm) from an LED. We imaged the surface of the DMD onto the surface of the 2D-programmable waveguide via a 4f set-up consisting of two tube lenses. The focal length of the tube lenses was chosen to demagnify the image of the DMD by a factor of 1.5, resulting in a projected pattern on the surface of the 2D-programmable waveguide with dimensions of 9.1 mm × 6.8 mm, with each individual pixel of the

projected pattern measuring 9 µm × 9 µm. Because the complex wave propagation spans a distance of 1 mm in the x direction, in practice, we used only a 9.1 mm × 1 mm region of the projected pattern to control the wave propagation in the 2D-programmable waveguide. Finally, although the DMD provides only binary modulation, we achieved continuous refractive index modulation by applying pulse-width modulation to the illumination pattern. This is feasible because the DMD can be switched on and off at a rate of 20 kHz, much faster than the resistance–capacitance time constant of the device, which is about 10 Hz.

To maximize the electro-optic effect in lithium niobate, we used high voltages of about 1 kV. We created sinusoidal voltages with an arbitrary function generator and amplified the voltage with a Trek 2220 high-voltage amplifier, which has a voltage gain of 200× and is capable of outputting voltages of up to 2 kV. We electrically contacted the device using high-voltage-rated probe arms with BeCu probe tips; one probe tip was put in contact with the gold electrode on top of the device, whereas a grounded probe tip touched the silicon substrate (Supplementary Fig. 9).

We used an a.c. frequency of 10 Hz for the experiments shown in Fig. 2 and 26 Hz for the experiments shown in Figs. 3 and 4. Because a.c. voltage was applied, the desired refractive index distribution was realized at the peak of the sinusoidal modulation. Thus, we triggered the camera to this peak, which explains the use of a higher frequency for the machine learning demonstrations to maximize the update rate of the experiment. Finally, we opted for a.c. over d.c. operation because the resistivity of the silicon dioxide cladding exceeds the dark resistance of the photoconductor; under d.c. operation, the photoconductor would not modulate the overall impedance of the circuit. Future modifications can enable d.c. operation, such as using an alternative cladding material that is more conductive or by increasing the dark resistivity of the photoconductor (Supplementary Information section 1B).

To measure the output of the computation performed by our device, we imaged the output facet of the device with an infrared camera (Allied Vision Goldeye CL-033). We built a 4f relay with a magnification factor of 5.3, allowing us to image the intensity distribution at the output facet, $I_{\mathrm{camera}}(x,y)$, with a resolution of 2.8 µm per pixel and a field of view of 1.7 mm in the x direction. We defined a small range of y to be the region of interest and integrated the intensity over this range to obtain the 1D intensity output of the 2D-programmable waveguide: $I_{\mathrm{out}}(x) = |E(x, z = L)|^2 = \int_{y_{\min}}^{y_{\max}} I_{\mathrm{camera}}(x,y)\mathrm{d}y$. We used an exposure time on the order of 500 µs, chosen to be much shorter than the period of the a.c. voltage, which was approximately 40 ms.

Data availability

All data generated during this work are available via Zenodo at https://doi.org/10.5281/zenodo.10775721 (ref. 51).

Code availability

An expandable demonstration code for simulating wave propagation through programmable waveguides is available at https://github.com/mcmahon-lab/2D-programmable-waveguide. The code used for this work is available via Zenodo at https://doi.org/10.5281/zenodo.10775721 (ref. 51).

References

49. Bender, N. et al. Depth-targeted energy delivery deep inside scattering media. *Nat. Phys.* **18**, 309–315 (2022).
50. Frumker, E. & Silberberg, Y. Phase and amplitude pulse shaping with two-dimensional phase-only spatial light modulators. *J. Opt. Soc. Am. B* **24**, 2940–2947 (2007).
51. Onodera, T. & Stein, M. Data repository of 'Arbitrary control over multimode wave propagation for machine learning'. *Zenodo* https://doi.org/10.5281/zenodo.10775721 (2025).

Acknowledgements

We acknowledge the US Air Force Office of Scientific Research for funding under award number FA9550-22-1-0378 and the National Science Foundation for funding under award number CCF-1918549. We thank NTT Research for their financial and technical support. This study was performed in part at the Cornell NanoScale Facility, a member of the National Nanotechnology Coordinated Infrastructure, which is supported by the National Science Foundation (grant number NNCI-2025233). P.L.M. acknowledges financial support from The David and Lucile Packard Foundation Fellowship. We acknowledge helpful discussions with C. Alpha, N. Bender, J. Clark, A. D'Addario, N. Flemens, J. Grazul, R. Hamerly, D. Heydari, P. Infante, O. Jaramillo, V. Kremenetski, M. Krenn, K. Li, G. McMurdy, R. Panepucci, C. Poitras, S. Prabhu, A. Windsor, F. Wu and Y. Zhao.

Author contributions

T.O., L.G.W. and P.L.M. conceived the project. M.M. Stein, T.O., L.G.W. and P.L.M. designed the devices and experiments. T.O., M.M. Stein, B.A.A. and R.Y. performed the device fabrication with aid and recipe development from M.R.S., M.B., M.J. and T.P.M. G.S. supervised M.R.S. and M.B. M.M. Stein, T.O., M.M. Sohoni and T.W. designed and built the imaging set-up to program the refractive index patterns. T.O. and M.M. Stein built the high-voltage and beam shaper set-ups, performed the experiments and analysed the results. T.O. derived the scaling law. M.M. Stein, T.O., L.G.W. and P.L.M. wrote the paper with input from all authors. P.L.M. supervised the project.

Competing interests

T.O., M.M. Stein, M.R.S., L.G.W. and P.L.M. are listed as inventors on a patent application (WO2023220401A1) on the 2D-programmable waveguide. The other authors declare no competing interests.

Additional information

Extended data is available for this paper at https://doi.org/10.1038/s41567-025-03094-2.

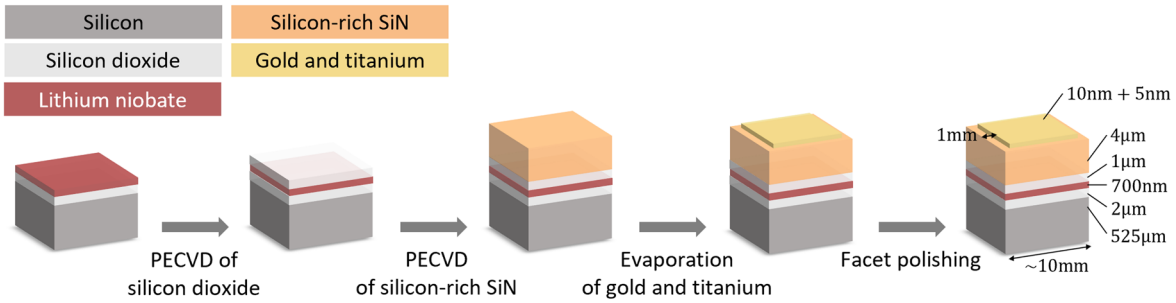
Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41567-025-03094-2.

Correspondence and requests for materials should be addressed to Martin M. Stein or Peter L. McMahon.

Peer review information *Nature Physics* thanks Daniel Brunner and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

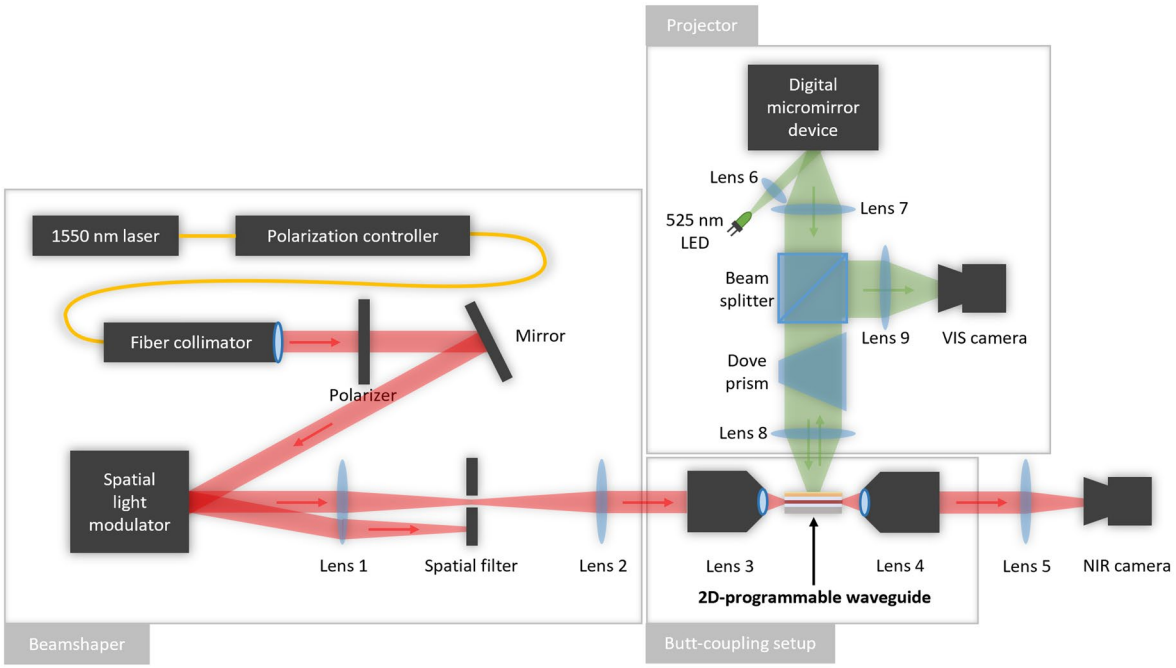
Reprints and permissions information is available at www.nature.com/reprints.

https://doi.org/10.1038/s41567-025-03094-2



Extended Data Fig. 1 | Fabrication process and device geometry of the 2D-programmable waveguide. Devices were fabricated from thin-film lithium niobate on insulator wafers consisting of a heavily doped p-type silicon substrate with a resistivity of 0.01 to 0.05$\Omega \cdot$ cm, a 2 μm PECVD silicon dioxide layer, and 700 nm Z-cut MgO-doped lithium niobate bonded by ion-cutting. After dicing, we deposited 1 μm of SiO_2 cladding and 4 μm of silicon-rich SiN (photoconductive layer) by PECVD with alternating plasma pulses to minimize stress. Electrodes were formed by evaporating 10 nm Ti and 5 nm Au, with tape masking to prevent breakdown at chip edges. Waveguide facets were polished using progressively finer abrasives to reduce coupling loss. The final device has facet dimensions of approximately 10 mm × 10 mm.

https://doi.org/10.1038/s41567-025-03094-2



Extended Data Fig. 2 | Schematic of the experimental setup. The setup consists of five key units: (1) a beamshaper to generate spatially varying one-dimensional optical inputs to the 2D-programmable waveguide, (2) a digital micromirror projector to impose a programmable illumination pattern controlling the refractive-index distribution, (3) butt-coupling optics to couple light in and out of the device, (4) (not shown) a high-voltage source applying oscillating fields across the electrode and conductive silicon wafer, and (5) an infrared camera to record the output intensity profile. The beamshaper provides full amplitude and phase control of input fields, while the projector enables spatial refractive-index modulation with continuous pixel-values via pulse-width modulation.