natureresearch
https://doi.org/10.1038/s41586-020-03070-1

Supplementary information

Parallel convolutional processing using an integrated photonic tensor core

In the format provided by the authors and unedited

Supplementary Materials for

Parallel convolution processing using an integrated photonic tensor core

J. Feldmann[1,], N. Youngblood[2,3,*], M. Karpov[4], H. Gehring[1], X. Li[2], M. Stappers[1], M. Le Gallo[5], X. Fu[4], A. Lukashchuk[4], A.S. Raja[4], J. Liu[4], C.D. Wright[6], A. Sebastian[5,#], T.J. Kippenberg[4#], W.H.P. Pernice[1,7,#] and H. Bhaskaran[2,#]*

[1] Institute of Physics, University of Muenster, Heisenbergstr. 11, 48149 Muenster, Germany
[2] Department of Materials, University of Oxford, Parks Road, OX1 3PH Oxford, UK
[3] Current Address: Department of Electrical and Computer Engineering, University of Pittsburgh, 3700 O'Hara St., Pittsburgh, PA 15261, USA
[4] Laboratory of Photonics and Quantum Measurements, Swiss Federal Institute of Technology Lausanne (EPFL), Station 3, CH-1015, Lausanne, Switzerland
[5] IBM Research - Zurich, Säumerstrasse 4, 8803 Rüschlikon, Switzerland
[6] Department of Engineering, University of Exeter, Exeter, EX4 4QF, UK
[7] Center for Soft Nanoscience, University of Münster, 48149 Münster, Germany

* These authors contributed equally.

#Correspondence to: wolfram.pernice@uni-muenster.de, harish.bhaskaran@materials.ox.ac.uk, kippenberg@epfl.ch, ASE@zurich.ibm.com

Table of Contents

1. Measurement setups

Single-vector convolution processing

The experimental setup used to carry out the convolution operations with a 9×4 matrix (see Figure 3 of the main paper) is shown in Fig. S1. Using a 40-mW pump laser amplified to approximately 1 W, a high-Q ring resonator made from silicon nitride is employed to generate a frequency comb in the single soliton state (see more details in the methods section).

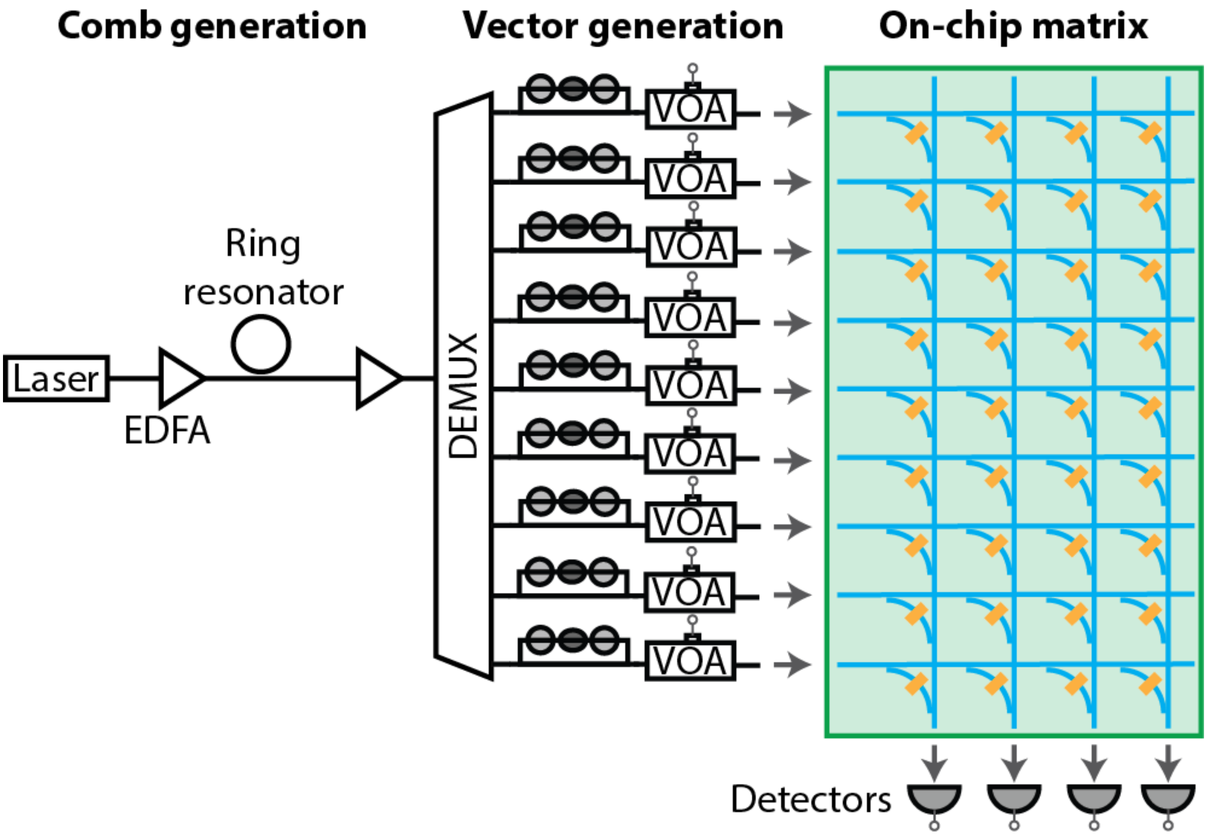

Fig. S1. Experimental setup for matrix-vector multiplication with a 9×4 matrix. A frequency comb is generated using a high-Q Ring resonator pumped by a CW laser. The input vectors are modulated using VOAs and send to the matrix input. The output signal is collected from the matrix columns using photodetectors.

After amplifying the comb lines with a further erbium-doped fiber amplifier (EDFA), the spectral lines are separated using a fiber based demultiplexer in the telecommunication window. The vector entries are encoded in the amplitude of the nine comb lines (Channels 27-35 of the ITU grid) and modulated using variable optical attenuators (VOAs). The electrical control of the VOAs is achieved using a custom-made digital-to-analog conversion (DAC) board interfaced via computer and an Arduino microcontroller. After propagation through the on-chip matrix and performing the multiply-accumulate operations, the output signal is recorded with low-noise photodetectors (New Focus Model 2011).

Using a matrix of size 9×4, four image kernels of size 3×3 are inscribed in the matrix (one kernel per column), resulting in 9×4=36 MAC operations per time step. The speed at which the convolutions can be performed in this particular setup is determined by the cutoff frequency of the modulators and the DAC-board, leading to approximately one operation per millisecond. The same setup is also used for the experiments with optical post-processing. In this case, the signal of the reference column was subtracted from the column with the filter-kernel.

Parallel vector convolution processing

In order to process several input vectors at the same time, the setup in Fig. S2 is used. 16 lines from the frequency comb (aligned to channels 27-42 of the ITU-grid) are separated using a fiber-based wavelengths de-multiplexer.

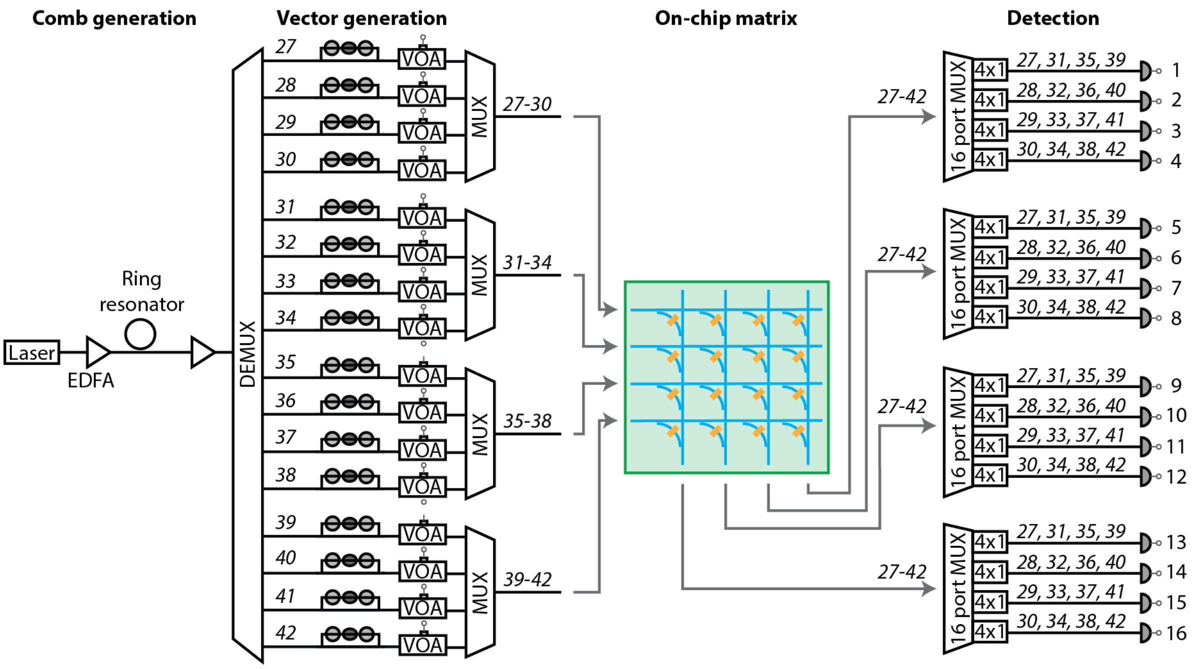

Fig. S2. Experimental setup for matrix-vector multiplication with a 4×4 matrix (multiplexing of four input vectors). Sixteen lines from the frequency comb are separated using a wavelength demultiplexer. After modulation with VOAs always four of the channels are grouped together and send to the matrix device. The signal of each output waveguide therefore contains all 16 wavelengths and is demultiplexed into the different wavelengths that represent a single vector. The numbers in italic denote the wavelength channel numbers according to the ITU grid.

The polarization in all channels is adjusted to optimize the coupling efficiency to the TE mode of the waveguides on chip. The amplitude of each channel is again modulated using variable optical attenuators. Each input vector is composed of four different wavelengths (for example 27, 31, 35, 39 for the signal going to detector 1). The wavelength multiplexers after the modulators combine all vector entries corresponding to a certain row and are sent to the matrix input.

At each of the matrix output columns all sixteen wavelengths are present, and have to be demultiplexed according to the required vector entries. Each column therefore gives four output values (for example detectors 1-4 for the last output column) leading to a total of sixteen full convolution operations in one step consisting of 64 multiply-accumulate operations.

Performance measurement

To find a lower boundary for the modulation speed, the input signals of a matrix with nine rows are modulated using the setup shown in Fig. S3a. For this experiment the same setup is used as for the 9x4 matrix, but in this case only one input of the matrix is modulated at a time using a fast electro-optical modulator (Optilab IMC-1550-20-PM). A vector network analyzer (Keysight ENA E5063A) is used to drive the modulator and to record the frequency response of the system.

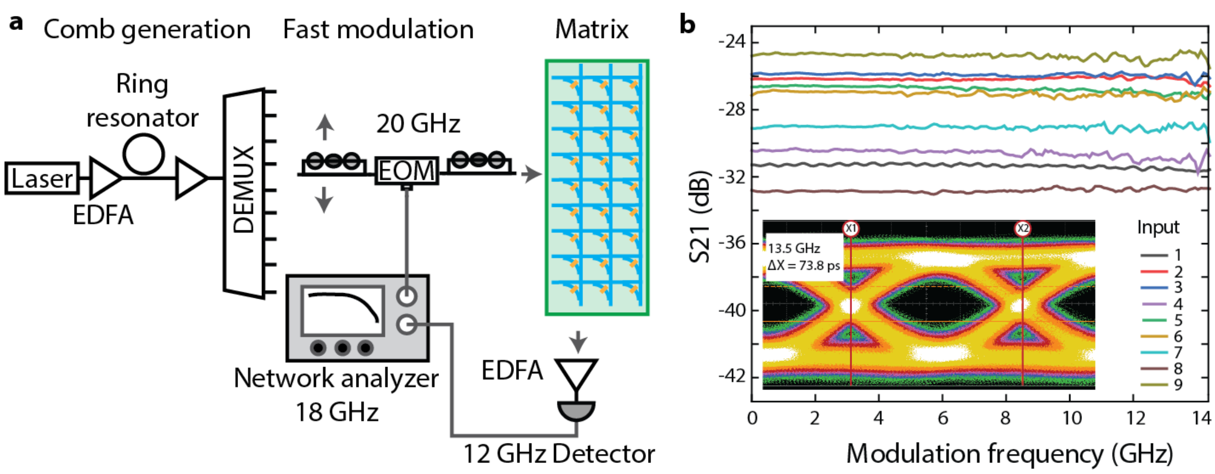

Fig. S3. Fast modulation of the input vectors. a) Each input of a 9×4 matrix is modulated individually and the frequency response is recorded with a network analyzer. **b)** Modulation bandwidth of all comb lines, showing modulation rates up to 14 GHz. Inset: open eye diagram at 13.5 GHz.

Fig. S3b shows the corresponding frequency response of the setup. Each output is recorded separately by measuring the whole system and subtracting a reference measurement characterizing only the off-chip components. The reference is obtained by sending the modulated signal after the EOM directly to the photodetector without passing through the on-chip matrix. The S21 parameter shows a flat response up to a frequency of 14 GHz, setting a lower boundary for the modulation speed of the input vectors. The speed is in this case limited by the slowest component in the setup which is the 12 GHz (3 dB bandwidth) Photodetector (New Focus Model 1554-B). The inset shows an open eye diagram at 13.5 GHz.

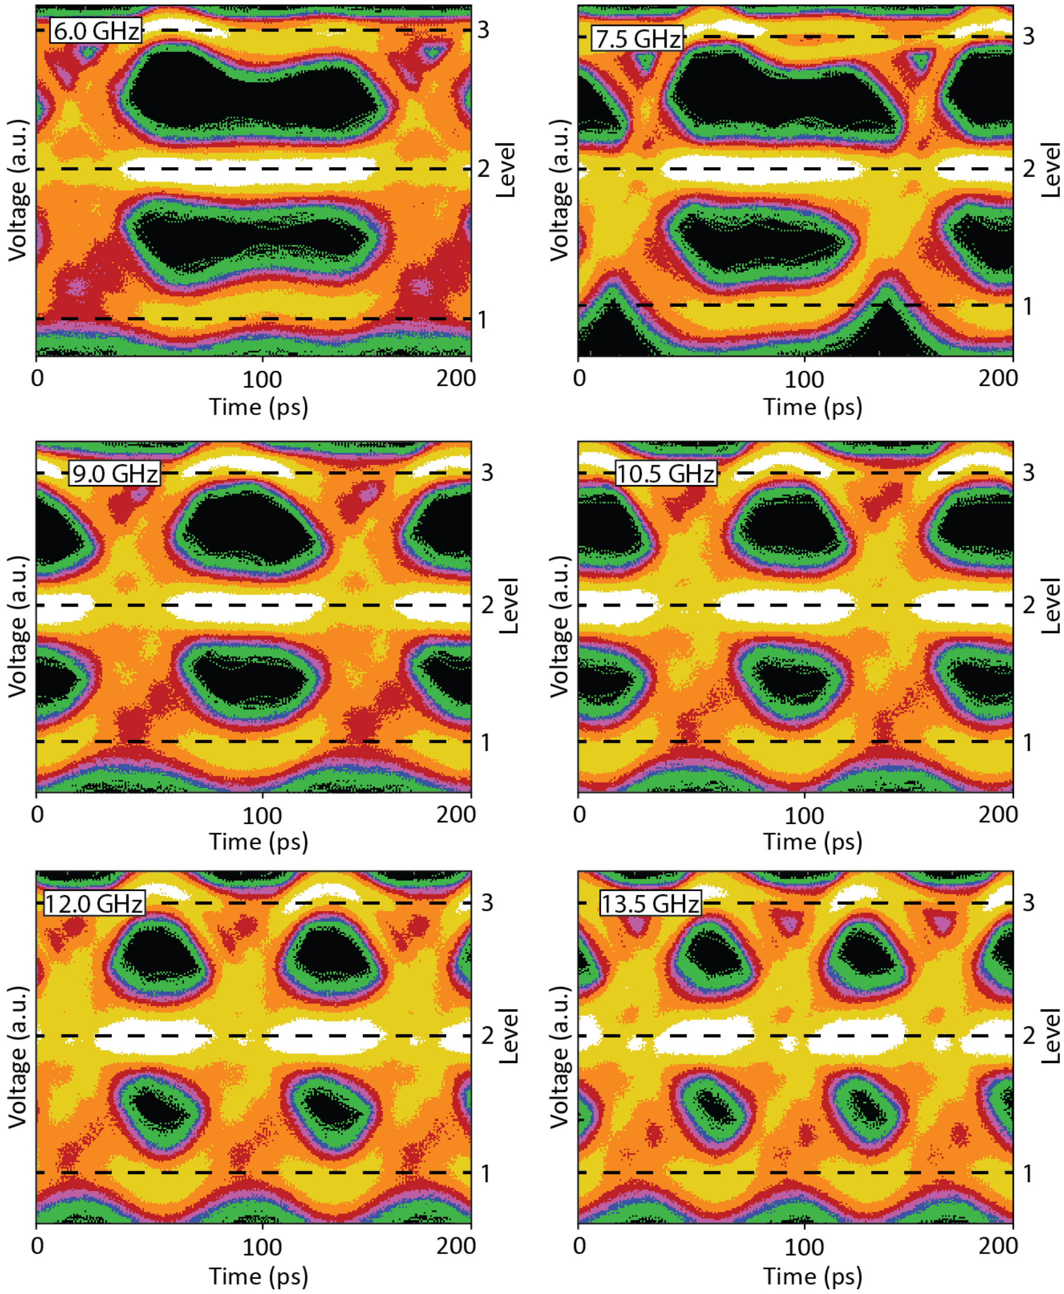

Fig. S4. Matrix-vector multiplication at different speeds. The input vector to a 2x1 matrix was modulated at speeds between 6.0 GHz and 13.5 GHz using PRBS patterns.

In the last experiments only a single input of the matrix was modulated at a time proofing that the modulation of the input vectors and the multiplication works with speeds up to at least 14 GHz. To show that also the incoherent addition of the different input rows can be performed at such high modulation frequencies, a full matrix multiplication with a 2x1 matrix is carried out.

Fig. S4 shows eye diagrams of the matrix output for different modulation speeds (6 GHz-13.5 GHz). The input patterns for both inputs were pseudo-random bit patterns (PRBS) leading to three different levels visible in the eye-diagram. Even at a modulation speed of 13.5GHz, all these levels can be clearly distinguished.

2. Setting the matrix elements

The matrix elements are assigned by controlling the phase composition of the PCM cells. The phase state is set all-optically using a separate input to each cell. Since the phase change is induced using a single wavelength, for convenience grating coupling devices are employed which are fabricated alongside the photonic circuitry and do not require an additional 3D DLW step.

Fig. S5 shows a scanning-electron micrograph of a single cell of the matrix device. The phase-change material (denoted PCM) serving as the matrix element is deposited on top of a waveguide crossing between the two directional couplers.

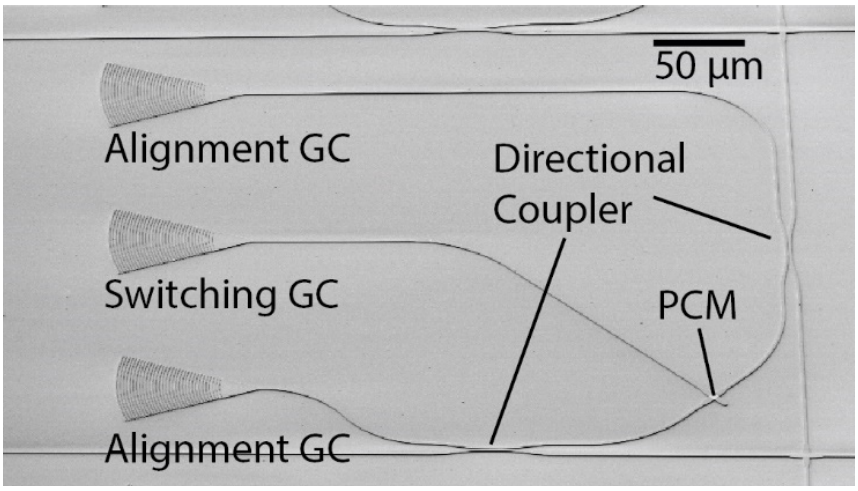

Fig. S5. Single matrix cell. Bragg grating couplers are used for switching the PCM-cells.

The two outer grating couplers (Alignment GCs) are used for aligning the fiber-array with respect to the photonic circuit in order to sufficiently couple light to the chip. This is achieved by maximizing the transmission of a probe signal between these ports. Switching of the PCM-cell is achieved using the middle coupler through the waveguide directly leading to the crossing with the PCM. Optical (rectangular) pulses with a width of 200 ns and different energies are employed to set the matrix element to the desired state. The change in transmission is observed through the alignment grating couplers.

The extra input port for switching is needed, because the directional couplers (DCs) in some of the matrix cells are required to have 100% transmission to the cross port of the DC. Therefore, these particular matrix cells would not be able to provide sufficient power levels for switching the PCM. Moreover, the power would also be transferred to adjacent PCM cells which are not intended to switch. In this case, to carry out the alignment through the two outer ports, the splitting ratios of the couplers were reduced to approximately 95% to be able to still detect the probe light.

Using single shot programming, distinguishing 34 different levels has already been demonstrated using optical switching [1]. To further increase the precision the matrix elements can be programmed with, an iterative process can be used as shown in Fig. S6.

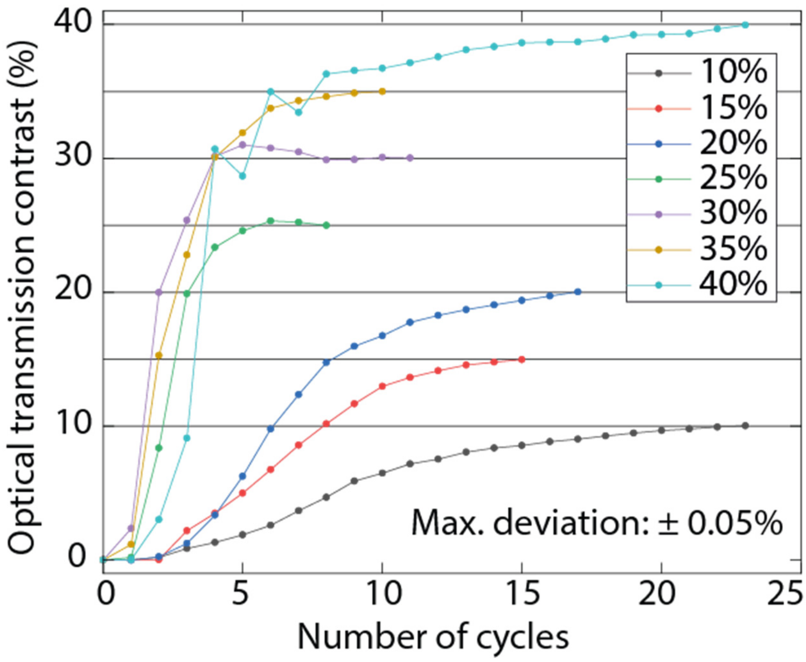

Fig. S6. Precision of programming matrix elements. Using an iterative process, the matrix elements can be programmed to arbitrary levels limited by the noise of the readout circuit.

By sending optical switching- and readout-pulses consecutively, arbitrary transmission states of the PCM-cells can be achieved limited only by the noise of the readout circuit. In Fig. S6 first a target transmission contrast (10%-40%) together with a maximum deviation of ±0.05%

from this value are defined. Then a switching pulse is sent, and the new transmission value recorded. Depending on the difference between the target transmission and the actual value the pulse energy for the next switching pulse was chosen. This process is then repeated until the optical contrast lies within the defined range of ±0.05%. This way in principle 400 matrix values can be used. In the given examples in Fig. S6 the number of cycles needed to achieve the wanted precision is up to 24 but can be optimized by more carefully calibrating the switching energies in every iteration. The achievable precision defined by the noise of the readout of the values and the control of the switching pulses.

3. Design of the DC splitting ratios used in the matrix

Figure 2a of the main paper shows a sketch of the matrix design indicating the splitting ratios of the different directional couplers that are chosen to provide an equal power splitting of the incoming light to the individual PCM-cells of the matrix. The horizontal directional couplers split the light to the different columns (image kernels) and the vertical directional couplers combine (addition) the light from the different rows after interaction with the PCM-patches (multiplication) within a certain column. Neglecting loss of the waveguide crossings, directional couplers and propagation, the splitting ratios for the vertical couplers can be calculated as $\frac{1}{j}$ with j being the number of the matrix row. The ratios for the horizontal couplers additionally depend on the number of columns n and are given by $\frac{1}{n+1-i}$ with the column index i.

In our photonic circuits, the experimentally determined waveguide crossing loss is approximately 0.2 dB and must therefore be considered in the design of larger matrices. Additionally, as explained in the previous section ("Setting the matrix elements") it is desirable to avoid splitting ratios of 100% for the directional couplers to be able to precisely control the matrix elements.

With the row index i and the column index j the transmission to the through port of a certain directional coupler including the insertion loss factor of the crossing l_c and the loss of the directional coupler l_{DC} (l_c=l_{DC}=1 meaning no loss) can be calculated as

$$t_{i+1} = \frac{1}{l_c l_{DC}(1-t_i)+1}$$

for the vertical couplers (t_i being the transmission of the i-th coupler) and

$$t_{j-1} = \frac{1}{l_c l_{DC}(1 - t_j) + 1}$$

for the horizontal directional couplers.

4. Design of the directional couplers

In order to achieve the correct power distribution in the matrix, the splitting ratio between the cross and the through port of several directional couplers was adjusted by varying the coupling length. We experimentally determine the splitting ratio as a function of the coupling lengths by measuring several fabricated calibration devices. Figure S7a shows the expected $\sin^2$ dependence of the splitting to the cross port at a wavelength of 1550 nm. Within lengths of 0 µm to 63 µm all possible splitting ratios can be achieved. The waveguide separation in the coupling area is 200 nm. Because the tapering section where the waveguides are brought closely together also remains constant for a coupling length of x=0, the splitting ratio does not start at 100% transmission to the through port, but rather has an offset.

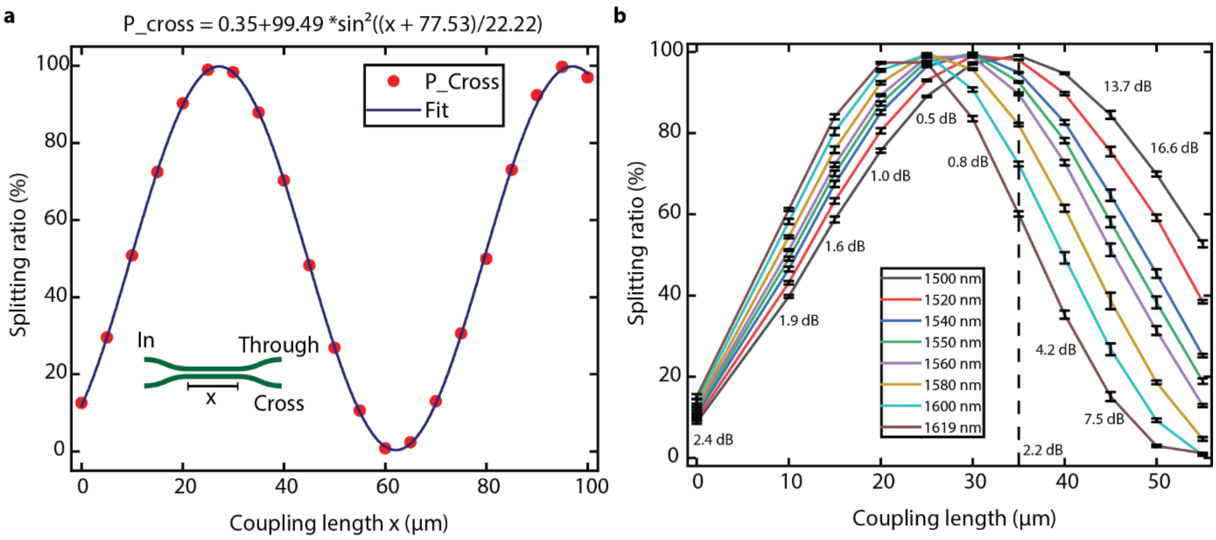

Fig. S7. Directional couplers. a) Splitting to the cross port of the directional couplers at 1550 nm. **b)** Wavelength dependence of the splitting ratio.

Figure S6b shows experimental data of the wavelength dependence of the directional couplers. The different lines depict the splitting ratio to the cross port of the coupler for wavelengths in the wavelength range from 1500 nm to 1619 nm for different coupling lengths. The numbers given in the plot show the maximum difference between two of the lines in dB and the vertical dashed line marks the position with the highest wavelength depended difference in the splitting ratio that was used in the matrix experiments in the main paper (2.2 dB). The results illustrate that a broad range of wavelengths can be used for multiplexing input vectors, that could even be increased by a less wavelength-dependent directional coupler design.

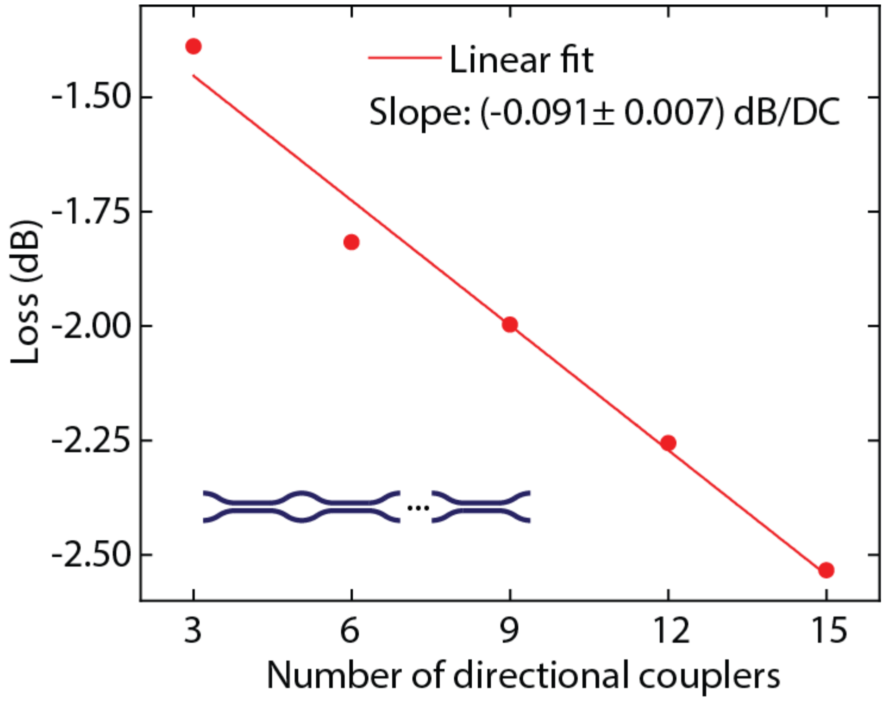

Fig. S8. Directional coupler loss. Optical loss introduced by a chain of directional couplers (DC). The slope gives the loss per coupler.

In order to estimate the loss of the directional couplers, several photonic devices consisting of chains of directional coupler as depicted in the inset of Fig. S8 were fabricated and the insertion loss measured as a function of the number of couplers (see Fig. S8). The slope of the linear fit reveals an insertion loss of approximately 0.1 dB/ coupler.

5. Precision of the photonic tensor core

Figure S9 shows the calculated standard deviation for a fixed-point scalar multiplication as a function of the precision of the input values. From the intersection with the horizontal line indicating the experimentally found standard deviation a resolution of 5 bits can be identified. The resolution is determined by the repeatability of the modulator settings, the noise of the detectors and the stability of the light source.

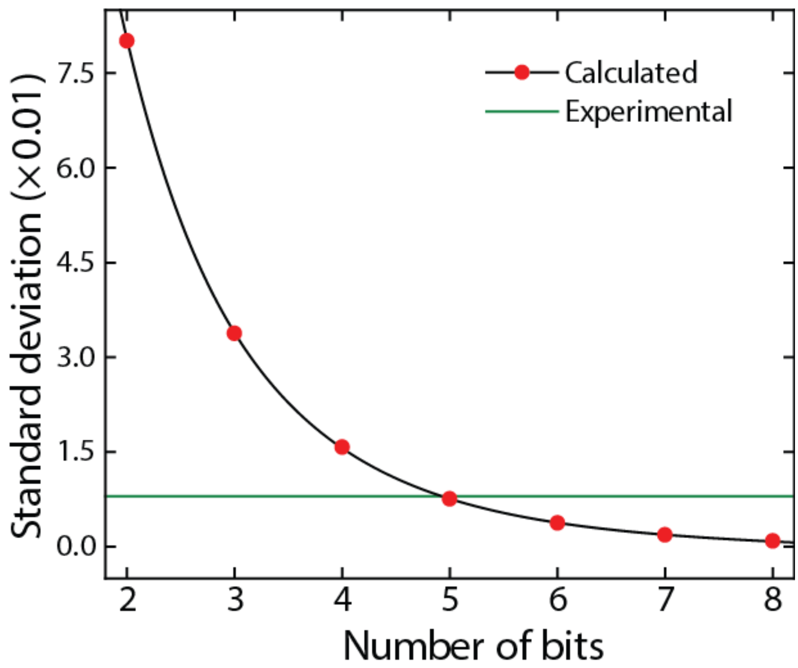

Fig. S9. Precision of the tensor core. Calculated expected standard deviation for fixed-point arithmetic with different resolutions. The green line indicates the experimentally obtained standard deviation for the photonic matrix and the intersection point therefore gives the resolution achieved with photonic MVM operations (5 bit).

6. Insertion loss of the photonic matrix

The loss of the photonic matrix limits its scalability and depends on three main factors: the loss of the directional couplers, the loss of the waveguide crossings and the loss due to the splitting of the light to the columns and incoherent combination of the rows. The propagation loss on the silicon nitride platform is below 1 dB/cm and can therefore be neglected. The contribution of the crossing and directional coupler loss scale linear with the matrix size m and the equal distribution of the light leads to a quadratic contribution to the loss resulting in the equation $A + B \cdot x + 10 \cdot \log_{10}(1/m^2)$. Figure S10a shows the measured optical loss for matrix sizes up to 32 x 32 and the corresponding fit, illustrating good agreement with the expected loss. Figure S10b shows an optical micrograph of an exemplary fabricated matrix of size 16x16.

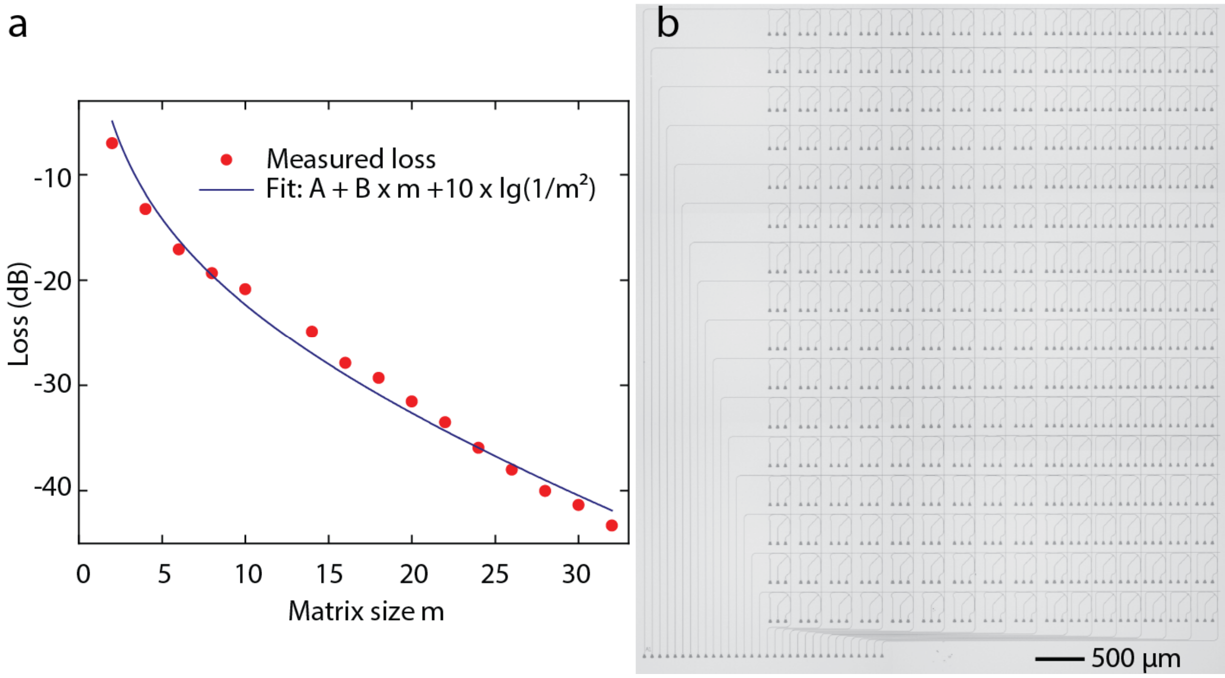

Fig. S10. Optical loss of the photonic matrix. a) Measured optical loss of fabricated matrices with different size up to 32x32. The fit confirms the expected dependence of the loss on the different loss factors inside the matrix. **b)** Optical micrograph of a fabricated matrix of size 16x16

7. Long-term stability

In order to investigate the long-term stability of the matrix-cells, the transmission through a waveguide with a PCM on top was measured for 600 hours (25 days), together with a reference measurement (see Fig. S11a). The photonic device used is shown in the inset of Fig. S11a. After crystallization of the PCM on a hotplate for 15 minutes at 230°C, it was optically switched by approximately 50%, before the measurement was started. The readout power, that stayed constantly on, was 0.5 mW (New Focus, Telecom Test Laser 6427) off-chip corresponding to approximately 20 µW inside both waveguides (reference and PCM). The reference transmission was optically maximized by moving the chip with a computer-controlled stage relative to the fibre array every 10 minutes. To reduce the influence of the optimization (optimized transmission for the reference port does not necessarily always correspond to the optimized transmission through the PCM), the realignment was carried out 5 times per datapoint and the average recorded. Both the curves for reference and PCM-transmission show a similar variation and stay constant on average of the time of the measurements. The periodic variations within the individual measurement reflect difficulties in the stabilization of the measurement setup over the time of the measurement, especially the coupling of the light to the chip using Bragg grating couplers and a fibre array.

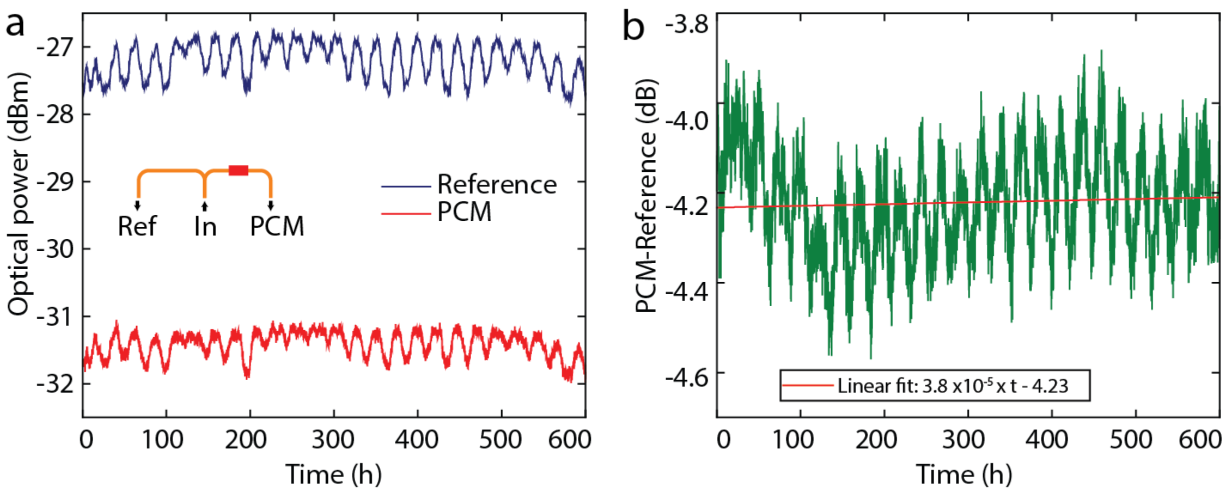

Fig. S11. Stability measurement. a) Transmitted optical power for a reference device (blue) and a PCM-cell (red) for 25 days. **b)** Transmission through the PCM-cell with the reference subtracted.

Fig. S11b) shows the PCM-transmission with the reference subtracted. Although the variations are not completely suppressed, the overall average given by the linear fit is almost constant, indicated by the small slope. It should be noted that a decreasing PCM-transmission is to be expected if the PCM recrystallizes as shown in [2]. The probe power to induce recrystallization needs to be ≥1 mW inside the waveguide on the silicon nitride platform.

Further data showing the long term stability of the optical transmission through the PCM-cell can be found in [3,4].

8. Detailed description of the matrix multiplication

To carry out the scalar multiplication $a \cdot b = c$, the first factor ($a \epsilon [0,1]$) is encoded in the amplitude of the optical signal ($a \cdot P_0$) and the second factor ($b \epsilon [-1,1]$) is inscribed in the phase state of a PCM-cell. The result c is obtained in a transmission measurement through a waveguide with the PCM-patch on top. As the optically transmitted power always depends on the input power, the measured optical output power does not directly reveal the product c. Considering for example the PCM-cell in the crystalline state corresponding to 0, the measured output power will be $P_1 = 1 \cdot P_0 \cdot l_{PCM}$ for the maximum input power ($a = 1$) and $P_2 = 0.5 \cdot P_0 \cdot l_{PCM}$ for half of the power ($a = 0.5$) with l_{PCM} being the loss factor of the phase-change material. Because the loss factor cannot be 0 (not all light is absorbed by the PCM, even in the crystalline state) $P_1 \neq P_2$. To correct for this offset and obtain the correct result of the multiplication, a reference value taken from the initial state of the PCM has to be subtracted from the measured output power, leading to the expression $c = \frac{P_{out} - P_{ref}}{\Delta P_{max}}$, with the measured output power P_{out} and the reference signal P_{ref}. The normalization factor ΔP_{max} is the maximum change in output power between the amorphous state and the reference signal: $\Delta P_{max} = P_0 \cdot l_{PCM}^{amorphous} - P_0 \cdot l_{PCM}^{ref}$. The measured output power is given by $P_{out} = a \cdot P_0 \cdot l_{PCM}$ and the reference power by $P_{ref} = a \cdot P_0 \cdot l_{PCM}^{ref}$. It is important to note that the reference power also depends on *a* and is not a fixed reference for all calculations.

Inserting P_{out} and P_{ref} into the expression for *c* leads to the equation

$$c = a \cdot \frac{\left(l_{PCM} - l_{PCM}^{ref}\right)}{\left(l_{PCM}^{amorphous} - l_{PCM}^{ref}\right)},$$

from which *b* can be identified as $b = \frac{\left(l_{PCM} - l_{PCM}^{ref}\right)}{\left(l_{PCM}^{amorphous} - l_{PCM}^{ref}\right)}$. Going from a scalar multiplication to a convolution including the sum of several products, P_{ref} also becomes a convolution that depends on the input vector. Therefore, for each full convolution operation of an image kernel with the input image, two individual convolutions (P_{out} *and* P_{ref}) have to be carried out, as experimentally shown in Fig. 3 of the main text.

In the experiments presented in the main text, the maximum contrast between the amorphous and crystalline state is approximately 4 dB. By setting the reference values (l_{PCM}^{ref}) to the center between the amorphous and crystalline PCM state ($l_{PCM}^{ref} = \left(l_{PCM}^{amorphous} - l_{PCM}^{crystalline}\right)/2$) also negative matrix entries can therefore be achieved.

9. Image kernels

In the main paper, several image filters are used to perform the convolution operations. For the example of a 9×2 matrix operated without post-processing in Fig. 3f an emboss filter of the form

$$\begin{pmatrix} -2 & -1 & 0 \\ -1 & 1 & 1 \\ 0 & 1 & 2 \end{pmatrix}$$

is employed.

The 2x2 kernels used when multiplexing four input vectors at the same time in Fig. 4 are

$$\begin{pmatrix}-1 & 1\\ -1 & 1\end{pmatrix}, \begin{pmatrix}-1 & -1\\ 1 & 1\end{pmatrix}, \begin{pmatrix}1 & 1\\ -1 & -1\end{pmatrix} \text{ and } \begin{pmatrix}1 & -1\\ 1 & -1\end{pmatrix}.$$

10.Energies and switching scheme for setting matrix elements

The pulses used for switching the matrix elements are of rectangular shape with a width of 200 ns. The maximum pulse energies used are approximately 800 pJ in order to achieve a contrast in the optical transmission of up to 70%. As the matrix elements representing the image kernels are usually fixed and do not have to be changed often, they were set manually by sending individual pulses with varying pulse energies until the desired state was reached with appropriate accuracy.

The switching process could be simplified using a double-step pulse scheme [5] leading to the final PCM state with a single pulse with high accuracy. Also, the pulse energies can be considerably reduced by using shorter pulse widths. Switching with pulses of a width of approximately a picosecond have been shown to yield switching energies of several pJ [6].

Because the matrix elements do not have to be changed often after the network is trained, the throughput and energy consumption per MAC (multiply-accumulate) operation is an important measure. In our experiments energy consumption is determined by the optical loss through the system. In order to characterize the loss of the complete setup, the optical power of a single wavelength before the first multiplexer (see Supplementary figure 1) and after passing through the matrix to the detector is measured, yielding approximately -45 dB for the 9×4 matrix and single vector operation and -51 dB for the 4×4 matrix and multiplexing of four vectors.

Considering the insertion loss of the individual components of the experimental setup (16-port MUX: -5 dB; Polarization controller and VOA: -1 dB; Matrix: -15 dB; GST: -3 dB; 2x coupling to/off chip: -9 dB), the measured overall system loss corresponds to the expected value. For the 4×4 matrix the loss is (by design) only -12 dB. In this case additional multiplexers are used at the output and the combination of the individual wavelengths is achieved using a 4×1 splitter giving additional loss of -6 dB and explaining the difference between the two experiments. It is important to note that the energy consumption can be significantly decreased by reducing the individual loss contributions in the setup. Especially the coupling loss can be improved to 1.5 dB per coupler (with optimized on-chip couplers) or even less when integrating the light source also on the same chip, thus gaining almost two orders of magnitude.

11. Compute density comparison and future scaling:

To compare computational throughput of our photonic tensor core with state-of-the-art digital and analog accelerators in both the optical and electrical domains (see Tables S1 for current optical and electrical implementations and Table S2 for future projections of optical approaches), we use the compute density metric TOPS/mm^2:

$$\text{Compute Density} = \frac{\text{Tera Operations Per Second (TOPS)}}{\text{Chip Area (mm}^2\text{)}}$$

In this metric, the photonic tensor core (SOI platform) outperforms any other hardware platform by at least one order of magnitude, including state-of-the-art electronic systems.

Technology	Clock Speed	Compute Density (TOPS/mm^2)	Efficiency (TOPS/W)	Precision (bits)	Ref.
Photonic WDM/PCM in-memory computing (SiN platform)	12 GHz[a]	1.2 (9×4 matrix)	0.4	5	This work
Photonic WDM/PCM in-memory computing (220nm SOI platform)	18 GHz	162[b] (8x8 matrix)	0.5	5	This work
Haswell E5-2699 v3 (CPU)	2.3 GHz	0.0039	0.018[c]	8	[10]
Google TPU (ASIC)	700 MHz	0.28	2.3[c]	8	[10]
Nvidia Tesla P40 (GPU)	1.3 GHz	0.10 0.025	0.19[c] 0.048[c]	8 16	[18]
TrueNorth (neuromorphic ASIC)	200 Hz	0.00027 (2.6×10^8 synapses)	0.8[c]	9	[19]
Neurogrid (neuromorphic ASIC)	3.7 Hz	0.0000026 (~8×10^9 synapses)	0.0021[c]	13	[20]
PCM crossbar array	4.2 MHz	3.6[d] (512×512 array)	28	8	[21]
HfO_2 memristor crossbar array	100 MHz[e]	0.066[e]	28	5–8	[22]

Table S1: Compute density and efficiency of current optical and electrical architectures.

a. Predicted speed based on high performance photonic integrated photodetectors, modulators, and high efficiency CMOS-based ADC. ADC power consumption[14] was 13mW/channel (6-bit, 12 GHz, Alphacore A6B12G) and 194mW/channel (8-bit, 25 GHz, Alphacore A8B40G).
b. See section S11 for more details.
c. Efficiency for Haswell, Google TPU, Nvidia Telsa P40, TrueNorth, and Neurogrid were estimated by dividing wall-plug power to operational TOPS rate per processor as described in [23].
d. Calculation projected based on a 14-nm fabrication node (rather than 90-nm node demonstrated).
e. Projected 100 MHz clock speed was deemed possible, but not experimentally demonstrated in this work. ADC power consumption[14] was estimated as 0.7mW/channel (10-bit, 320 MSPS, Alphacore A10B320M).

However, by upscaling the demonstrated prototype systems, substantial performance gains become possible, as collected in table S2.

Technology	Clock Speed	Compute Density (TOPS/mm^2)	Efficiency (TOPS/W)	Precision (bits)	Ref.
Photonic WDM/PCM in-memory computing (220nm SOI platform)	12 GHz[a] 25 GHz[a]	420[b] 880[b] (64×64 matrix)	7.0[c] 5.1[c]	5 5	Future scaling
Photonic coherent DNN (MZI)	25 GHz[a]	0.52 (4×4 matrix)	0.33[d]	~3.5	Future scaling[24]
Photonic hybrid laser NN (III-V/Si platform)	12 GHz	9.5 (56×56 matrix)	3.2[d]	5.1	Future scaling[12,25]

Table S2: Compute density and efficiency of projected optical architectures.

a. Predicted speed based on high performance photonic integrated photodetectors, modulators, and high efficiency CMOS-based ADC. ADC power consumption[14] was 13mW/channel (6-bit, 12 GHz, Alphacore A6B12G) and 194mW/channel (8-bit, 25 GHz, Alphacore A8B40G).
b. Calculation based on a 30 × 30 µm^2 photonic MAC unit cell with a nominal 5-µm bend radius.
c. Limit in the case of negligible insertion losses for couplers and crossbars ($l_{DC} \approx l_C \approx 1$).
d. Efficiency calculation includes average power consumption of 10 mW/heater and 5 mW/heater for refs[24,25] respectively.

For our SiN on insulator platform, the minimum dimensions of the unit cell is limited by the bend radius of the partially etched SiN rib waveguides, the waveguide crossing structures, and the additional grating couplers used to optically switch the GST. The footprint of this design can be significantly optimized by using a silicon-on-insulator (SOI) platform with a standard 220 nm × 500 nm strip waveguide design. Figure S12 shows an optical micrograph of an 8×8 matrix fabricated on SOI by a foundry (Cornerstone) with a matrix cell size of 127 µm x 127 µm (compared to 285 µm × 354 µm on silicon nitride as shown in the main paper).

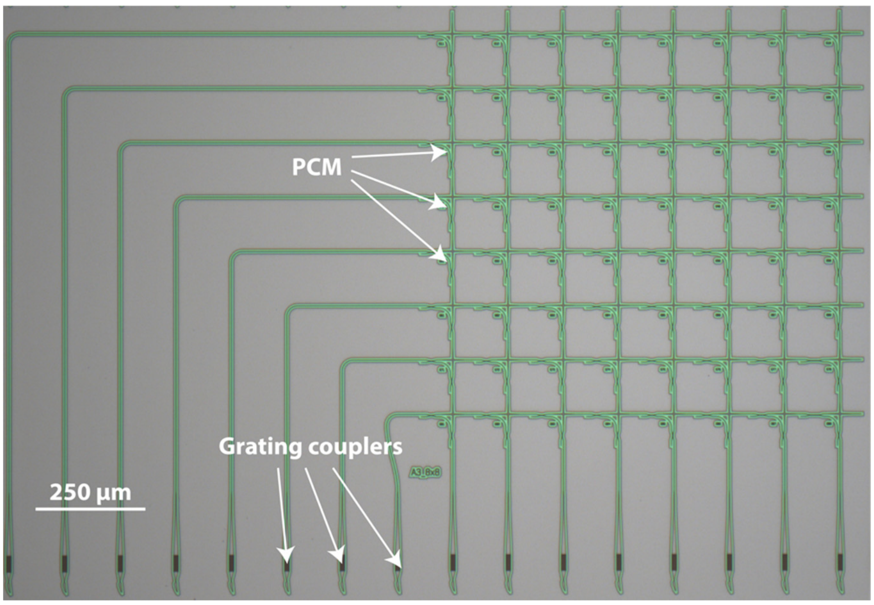

Fig. S12. Optical micrograph of an 8×8 matrix fabricated using silicon on insulator.

In order to estimate the spectral width of the frequency comb that can be exploited for computation, the optical transmission spectrum of the comb was measured in a static experiment after transmission through the on-chip SOI matrix (see Fig. S13a). Comb lines are visible in a range of more than 50 nm, limited by the spectral width of the input grating couplers used to couple the light onto the chip. With a line spacing of 100 GHz this leads to about 65 lines that can be used for parallel vector processing. Figure S13b additionally shows the frequency response of the matrix for modulation speeds up to 18 GHz (limited by the speed of the network analyser Keysight ENA E5063A), which leads to a compute density of 162 TOPS/mm² in the experiment on SOI.

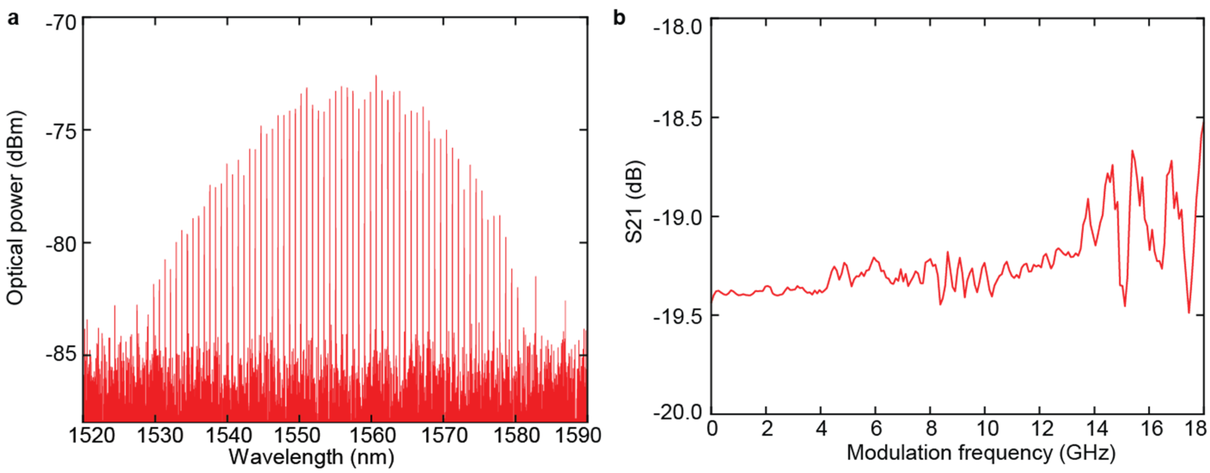

Fig. S13. SOI matrix of size 8×8. a) Optical spectrum of the frequency comb after transmission through the on-chip matrix. A range of approximately 50 nm can be exploited for computation. **b)** Frequency response of the on-chip matrix as a function of the modulation frequency, showing that modulation speeds up to 18 GHz can be used.

However, the footprint of the matrix used in the experiment was mainly limited by the size of the components (especially the waveguide crossings) available in the foundry. The footprint of the unit cell can therefore further be decreased to 30 × 30 μm^2 by optimizing the individual components. Here we use a compact (9 × 9 μm^2), low loss waveguide crossing design (–0.028 dB insertion loss for telecom band[7]) and 3dB directional coupler with a 100 nm waveguide spacing, 400 nm wide strip waveguide, and 10 μm coupling length[8] . We use a 3dB coupler to calculate the maximum MAC unit cell based on the coupling equations for our photonic crossbar array (**Supplementary Information, section 3**). The longest directional coupler length required for the array will be the 3 dB directional coupler (i.e. $t_{N-1} = 1/2$), since the transmission required for the coupler in the last column and first row will be $t_N = 1$ which can be accomplished with a simple connecting waveguide with a 5 μm bending radius.

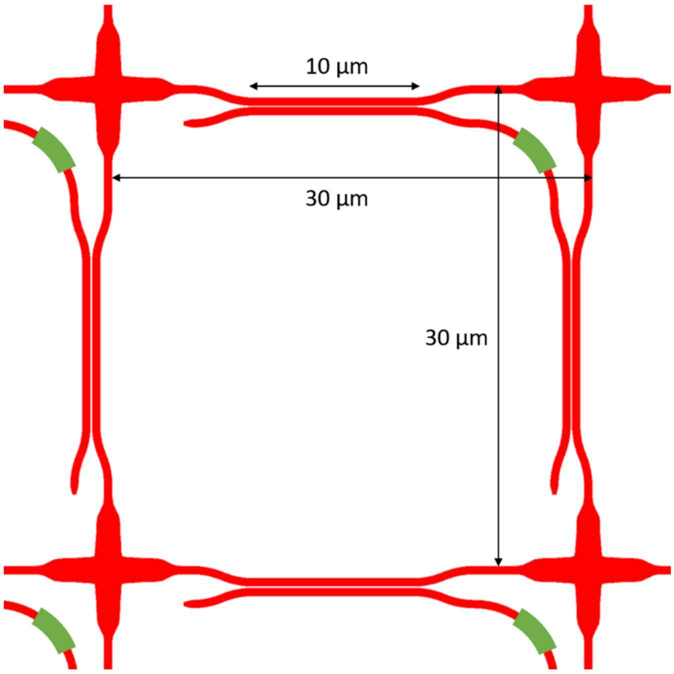

Fig. S14. CAD drawing of photonic MAC unit cell on a fully-etched 220 nm SOI platform. Compact waveguide crossing design and 3dB directional couplers based on [7,8].

Figure S14 shows our SOI unit cell with 3dB directional couplers connected by compact waveguide crossings. To control the state of the GST, compact waveguide-integrated heaters which have been recently demonstrated[9] would enable a scalable method for electrically programming the optical transmission of the GST (green patches in **Figure S14**) and replace the additional optical grating couplers and crossing (**Supplementary Information, section 2**). The footprint of this photonic MAC unit is 30 × 30 μm^2.

Assuming a modulation and detection speed of 25 GHz, together with a 16-channel multiplexing scheme, we calculate the compute density of our SOI photonic tensor core to be 880 TOPS/mm^2, which is 3 orders of magnitude greater than state-of-the-art digital approaches[10]. This shows the significant promise of our in-memory photonic approach which allows wavelength multiplexing to achieve parallel processing in a single photonic core without sacrificing speed or latency. Additional density may be gained by not restricting the design to a rectangular grid (i.e. reducing the unused space in the center of the unit cell).

To estimate the optical energy needed per MAC operation, we assume insertion losses of 0.028 dB per crossing[7], 0.1 dB per directional coupler, and 0.6 dB per GST/integrated electro-thermal switch[9]. Using an on-chip optical source and photodetectors to eliminate coupling losses, the transmitted power from the first input row ($i = 1$) to the first output column ($j = 1$) will be:

$$T_{1\to 1} = l_{GST} l_{DC}^{N} l_{C}^{N-1} (1 - t_{j=1}) \prod_{i=2}^{N} t_i$$

where l_{DC}, l_C, t_i and t_j are as defined in section 3. For an $N \times N$ matrix (with $N = 64$ as demonstrated in[11]), we have $t_{j=1} = 0.9947$, $t_{i=1} = 0$, $l_{GST} = 10^{-0.06}$, $l_{DC} = 10^{-0.01}$, and $l_C = 10^{-0.0028}$ resulting in a total transmission of $T_{1\to 1} = -46$ dB for amorphous GST. Due to the design of the crossbar array, any $T_{i\to j}$ should also equal $T_{1\to 1}$.

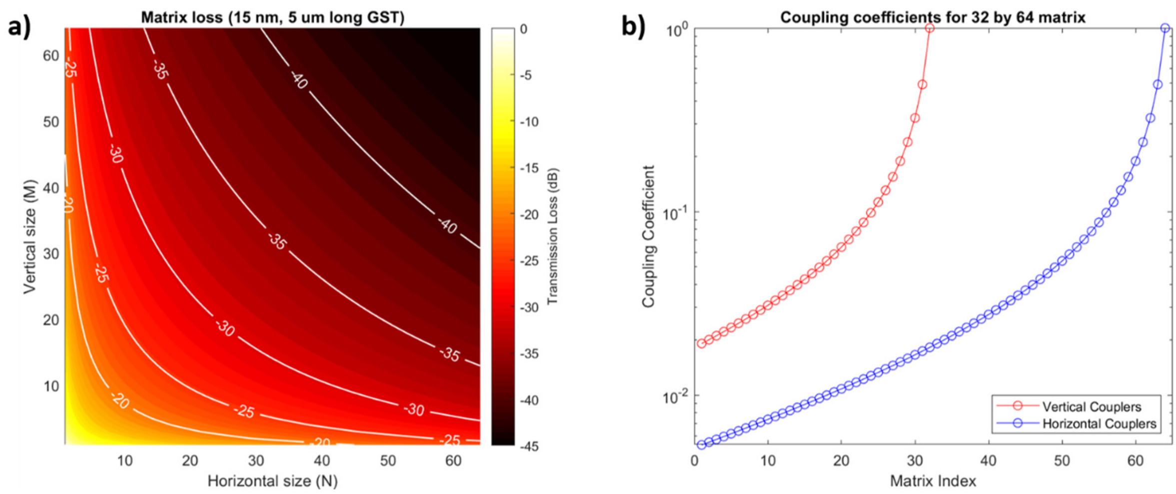

Fig. S15. Insertion loss and coupling coefficient calculations. (a) Overall insertion loss of an M×N photonic matrix. (b) Example values for coupling coefficients in a 32×64 array, which scale as 1/M and 1/N.

Figure S15 shows the calculated insertion loss as a function of matrix dimensions as well as an exemplary calculation of the coupling coefficient values for a 32×64 photonic matrix using the same insertion loss values as above.

The total energy consumption of the system to overcome shot noise and the capacitance of the photodetectors at the output with a fixed precision of N_b-bits of fixed precision[12] can be calculated from the following equation:

$$P_{total} = N \left(\frac{1}{N} \frac{h\nu}{\eta} \max\left(2^{2N_b+1}, \frac{C_d V_r}{e}\right) f_{mod} + P_{rec} + E_{mod} N_b f_{mod} \right) \times (\text{\# of WDM channels})$$

where $N = 64$, $h\nu$ is the photon energy for a center wavelength of 1.55 μm, $C_d = 2.4$fF and $V_r = 1$V are the capacitance and driving voltage of the photodetectors[13], and η is the combined quantum efficiency of the detector, laser, and optical loss through the crossbar array ($\eta \approx -56$ dB assuming a 10% combined quantum efficiency of the laser and photodetector). We estimate the energy consumption at the receivers and modulators to be dominated by the ADC[14] and optical modulators[12,15] ($P_{rec} \approx 200$ mW/channel and $E_{mod} \approx 1$ pJ/b). This gives a total power consumption of 639 W for a 64×64 array with 16 WDM channels or 5.1 TOPS/W. We compare the compute efficiencies of different architectures in Table S1. The future scaling potential of the presented photonic tensor core (as derived in this section) compared to other optical as well as electronic systems is estimated in Table S2.

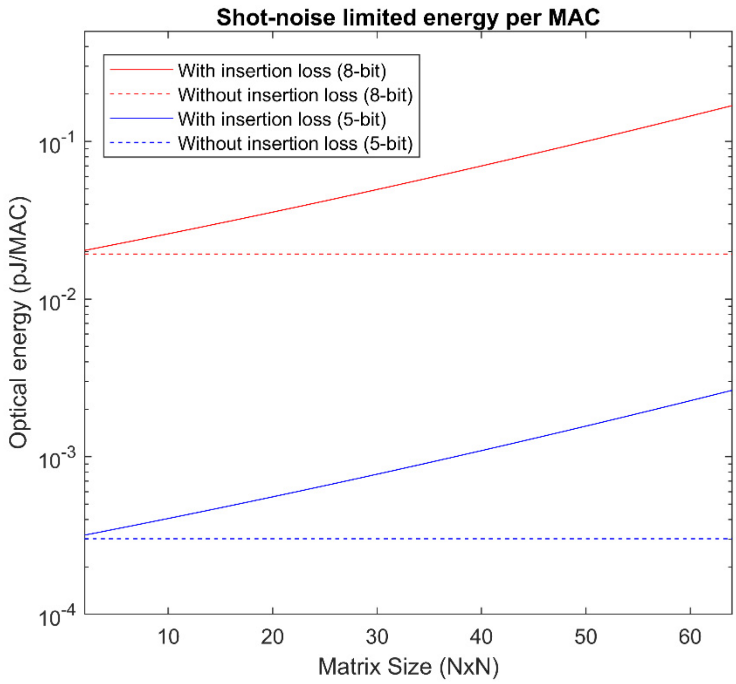

Fig. S16. Plot of optical energy per MAC required to overcome shot noise for 5-bit and 8-bit precision at the output. The dashed lines place a lower limit on the energy per MAC for a lossless crossbar array.

To better understand the energy scaling of the array itself without contributions from the modulation and detection input-output layers, we now consider the minimum optical power required to overcome the quantum shot noise for a fixed precision number at the output. We can compute the minimum optical energy per MAC using the following equation[12]:

$$E_{MAC} \geq \frac{1}{N^2} \frac{h\nu}{\eta_{array}} 2^{2N_b+1}$$

where η_{array} is the optical loss through the crossbar array. For the 64×64 matrix above with N_b = 8-bits of precision and $\eta_{array} = -46$dB, this corresponds to $E_{MAC} = 0.16$ pJ/MAC. Minimizing the insertion loss of the directional couplers and crossbar can lower $T_{i \to j}$ to a minimum of –36 dB which reduces the energy consumption to E_{MAC} = 17 fJ/MAC. The energy per MAC can be further reduced by reducing the precision at the output. Figure S16 plots the dependence of shot-noise limited energy per MAC as a function of photonic crossbar dimension and output precision. In the absence of insertion loss from the various components, E_{MAC} is a constant since $\eta_{array} \propto N^2$.

12. Broadband fiber-to-chip couplers

In order to gain access to the broad frequency comb spectrum we combine the broadband transparent nanophotonic waveguides with a broadband coupling scheme for out-of-plane chip access. We employed 3D-polymer structures[16] for interfacing planar circuitry with the optical fiber array used to perform the measurements. A schematic rendering of the 3D coupler is shown in Fig. S17a. The device consists of an adiabatic mode converter between planar waveguides and 3D-polymer-waveguides, a tapering region, in which the size of the mode is increased, a total reflection plane at which the beam is reflected out of the plane of the chip towards the fiber array and finally a lens for focussing the beam onto the fiber facet. We optimized the geometry of the coupler for operation in the telecom wavelength regime. The couplers were written using a Nanoscribe Professional System after finishing planar nanofabrication. We employed IP-Dip photoresist (Nanoscribe GmbH) and a slicing and hatching distance of both 100nm. The polymer couplers were connected to the nanophotonic waveguides after automatically aligning the structures to the markers shown in Fig. S15b using computer vision.

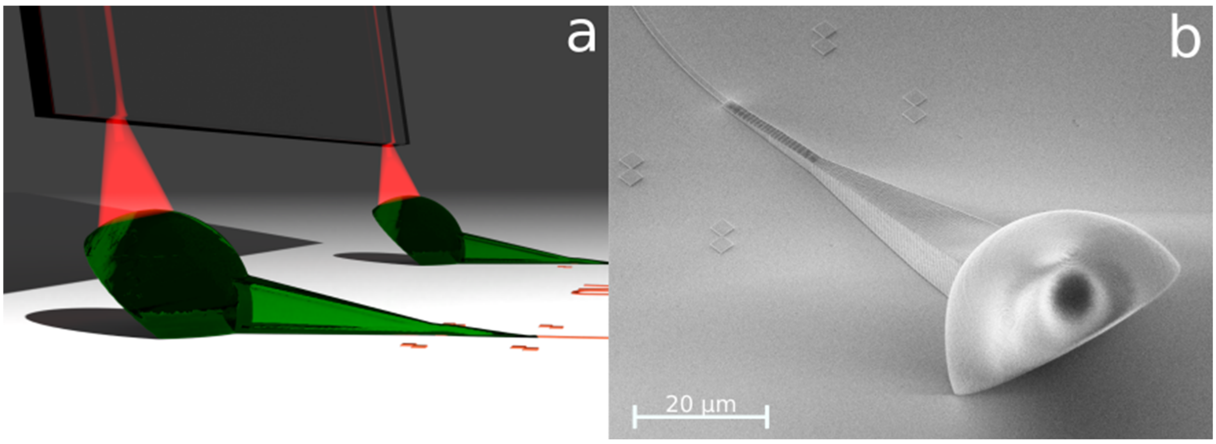

Fig. S17. Images of broadband 3D-couplers. (a) Rendering of two total internal reflection couplers, which interface planar circuitry with an optical fiber array. (b) SEM-image of a fabricated total internal reflection coupler.

The coupling structure was initially optimized using 2D numerical simulations as shown in Fig S18b in order to match the outcoupled beam to the mode of an optical fiber. Subsequently the design was further optimized experimentally to maximize the coupling efficiency by scanning the geometric parameters of the design. In order to characterize the transmission of the couplers, we characterized pairs of couplers which were connected by a planar waveguide on chip. A typical transmission spectrum is shown in Fig. S18a.

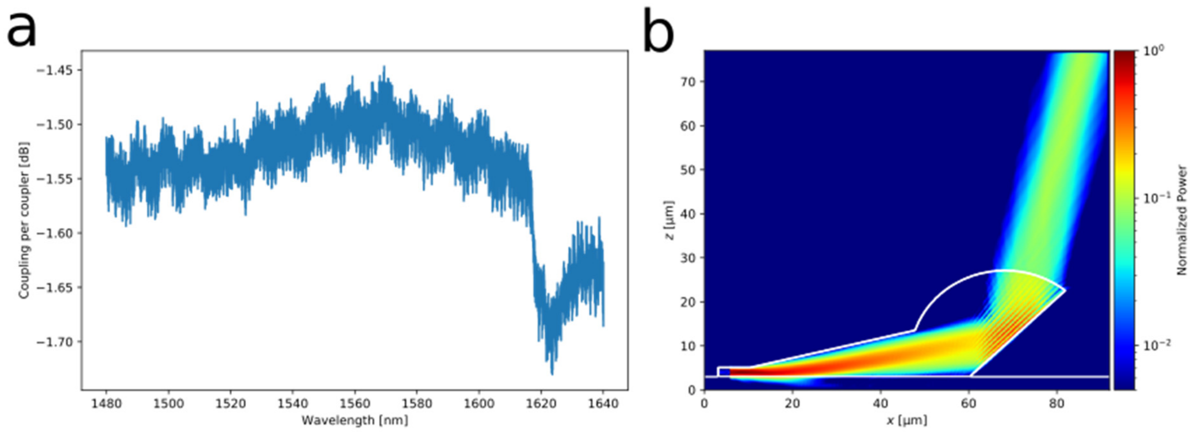

Fig. S18. Measurement and simulation of broadband 3D-couplers. (a) Measurement of the coupling efficiency of a total reflection coupler in the telecom wavelength range. (b) 2D-Simulation of a total reflection coupler using the FDTD-method [17].

By assuming equal coupling efficiency, we obtain the loss per coupler by dividing the measured transmission loss. For the optimized design, we obtained insertion loss below -1.5 dB. Importantly, the 3D couplers show very low wavelength dependence. In the measured window from 1480 nm to 1640 nm the peak transmission varies by less than 0.25 db. Thus, the 3D couplers provide efficient and broadband transmission to our photonic circuitry.

13. Broadband transmission of the frequency comb through the matrix

To estimate the number of comb lines that can be used for encoding input vectors, the spectrum of the frequency comb is measured after transmission through the on-chip matrix. Due to the use of broadband couplers described before, the transmitted spectrum in the silicon nitride platform (see Fig. S19) shows lines in a range of 150 nm from 1500 nm to 1650 nm. The inset shows the section of the spectrum around 1550 nm that was used for the experiments in the main paper.

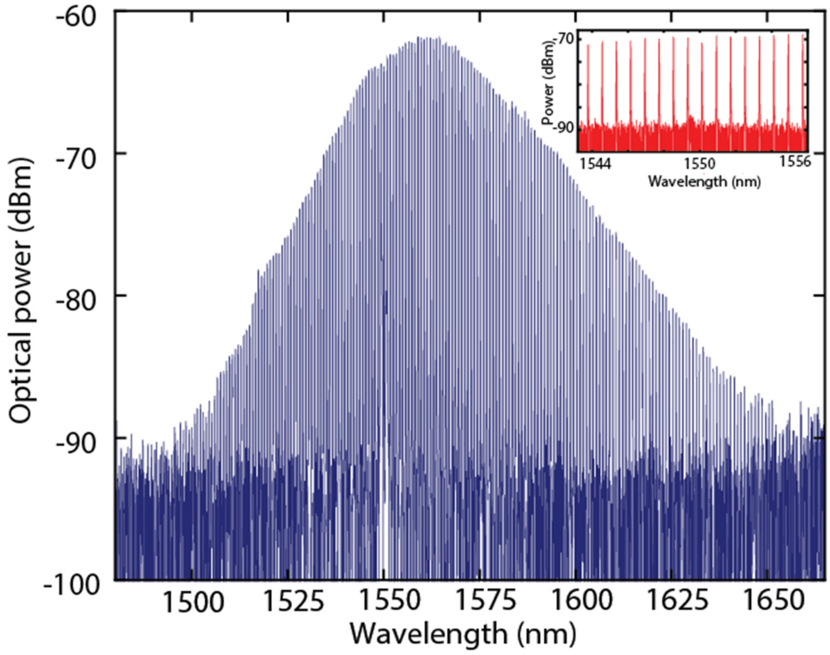

Fig. S19. Broadband operation of the tensor core. Transmission spectrum of the coherent soliton frequency comb through the device. Inset: spectrum of the 16 selected comb lines used in the experiment.

14. Convolutional neural network for MNIST digit recognition using the photonic tensor core

The following sections describe the design of the convolutional neural network used for digit recognition, and its experimental demonstration. The network was trained in software and the convolution kernels later programmed to the photonic tensor core executing the convolution operations experimentally on the test dataset. The electronic control was achieved using an FPGA running at 4 GS/s.

14.1 Design and training of the convolutional neural network for digit recognition

The convolutional network used for the experiment of digit recognition with the MNIST database was designed to suit the limitations of the experimental setup to using small matrices and demonstrating the potential of the photonic tensor core for accelerating neural networks. The network was designed and trained using the python interface keras with the tensorflow backend and consists of an input layer, a 2D convolution layer (valid padding) with four kernels of size 2x2 and a subsequent ReLU activation function. After flattening the outputs, a 10 neuron fully-connected layer with softmax classification is added.

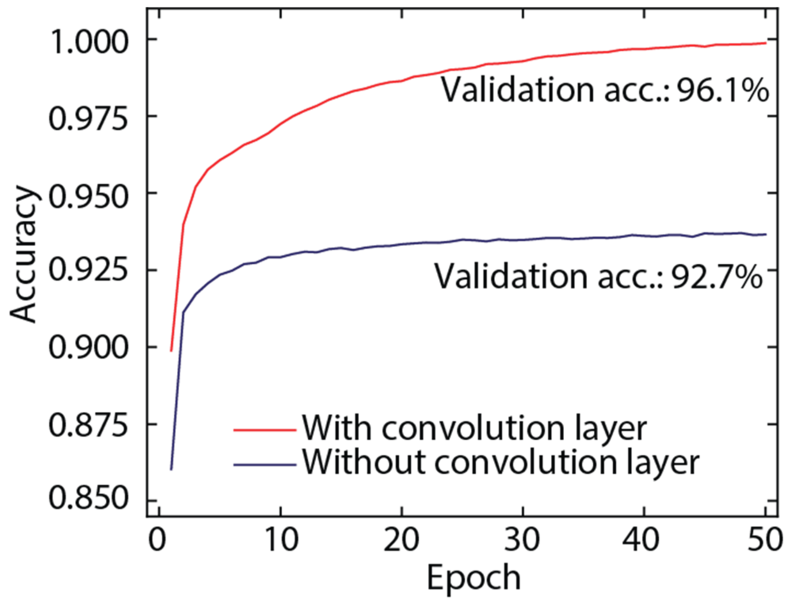

Fig. S20. Training of the convolutional neural network. History of the training of the network over 50 epochs. With a convolution layer the accuracy on the test-dataset is 96.1%.

The complete network was trained in software (see Fig. S20) leading to an accuracy on the test-dataset of 96.1%. As a comparison, the same network without the convolution layer reached a validation accuracy of 92.7%.

14.2 FPGA for high data-rate signal generation and data acquisition

In order to show convolution operations with sample rates in the GHz-regime, we employ a state-of-art FPGA with RF-DACs/ADCs (Zynq UltraScale+ RFSoC XCZU28DR-2FFVG1517E). Limited by the speed of the ADCs we set the sampling rate of the DACs/ADCs to 4 Gsamples/s. To achieve the necessary data rates, we utilize separate FIFO-BRAM-Buffers for the DACs and ADCs. The FPGA allows us to control 4 electro-optical modulators and read in data from a fast photodetector in parallel.

For synchronization, we block the data connection between the Buffers and converters until the Input Buffer (connected to the DAC) is filled with data from the DDR4-RAM using direct memory access (DMA). Subsequently, both data connections are unblocked, and the data is converted to analog data, send through the photonic convolution circuit, digitized by the ADC and written into the Buffer. The resulting data is written back into the DDR4-Memory via DMA.

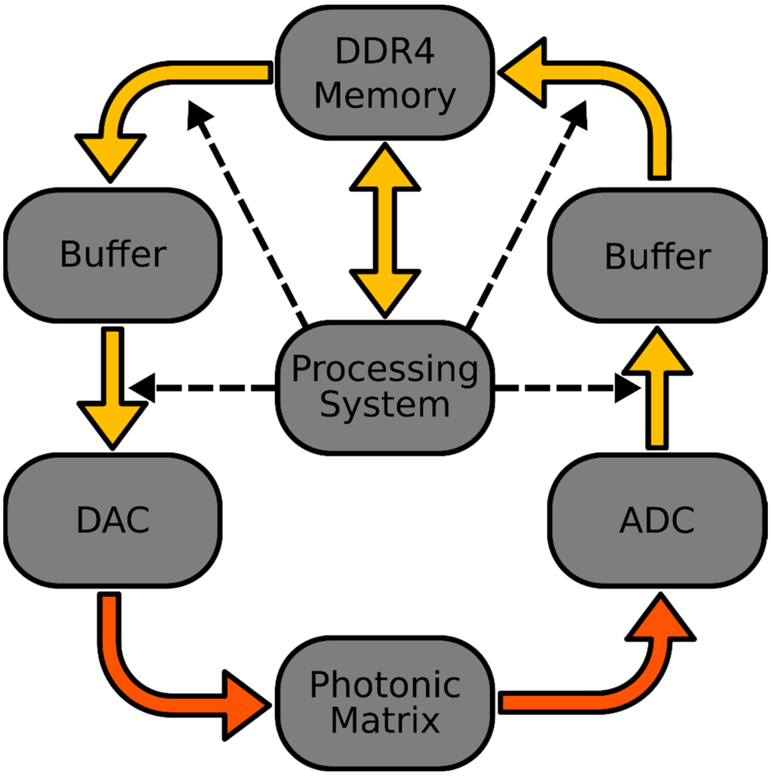

Fig. S21. Schematic representation of the high-speed signal generation and data acquisition using an FPGA. The loaded hardware layout consists of two FIFO-Buffers, a digital-to-analog converter (DAC) and an analog-to-digital converter (ADC). Both buffers are connected to the DDR4-Memory. Digital data connections are depicted in yellow, analog connections in red and control connections in dashed black. The processing system is connected to the DDR4-Memory to handle the data, to the data connections between DDR4-Memory and Buffers to load data into these/read it back and to the data connections between Buffers and DACs/ADCs to synchronize those.

14.3 Experimental methods

In the experiments implementing the convolutional neural network described in section S14.1 of the supplementary information (and Fig. 5 of the main text), the four 2x2 kernels of the convolution layer were processed sequentially in a 4x1 matrix. In a first step the 10000 input images of the test-set of the MNIST handwritten database were arranged into the input vector matrix as described in the main text. Subsequently, a reference measurement was taken (see section S8) by modulating the input vectors on the amplitudes of the four optical input signals to the photonic matrix using electro-optic modulators driven by the DACs of the FPGA described in the previous section. The output signal after passing the on-chip kernel (being in the reference state) was optically amplified and collected from a 12 GHz photodetector with an ADC of the FPGA. The DACs itself were operated at 4 GHz but every vector was applied for two clock cycles so that effectively a processing rate of 2 GHz is achieved. This is to account for the fact that the optical path lengths of the four signals in the off-chip setup did not exactly match, so that the modulated signals are slightly shifted in time. Delays down to the time of a single clock cycle (250 ps) could be regulated adding in electronic delays using the FPGA. To reduce the influence of noise, the reference measurement was repeated 100 times and the average calculated.

After measuring the reference, the images were processed with the kernel matrix column by column (again with averaging 100 times) and the reference values were subtracted. After normalization (see section S8) the experimentally processed convolution results were fed to a computer and the subsequent ReLU activation and fully connected layer were processed electronically, leading to a prediction accuracy of 95.3% compared to 96.1% achieved when processing the whole network with a computer.

14.4 Confusion matrices

Calculated

Prediction \ Label	0	1	2	3	4	5	6	7	8	9
0	98.8	0.00	0.48	0.00	0.10	0.34	0.84	0.00	0.72	0.50
1	0.00	98.9	0.68	0.00	0.00	0.00	0.31	0.58	0.00	0.69
2	0.10	0.26	93.7	0.59	0.31	0.00	0.10	2.04	0.51	0.20
3	0.00	0.26	1.45	96.2	0.00	1.12	0.00	0.58	1.03	0.59
4	0.00	0.00	0.58	0.00	97.2	0.00	0.21	0.19	0.51	1.49
5	0.20	0.09	0.00	1.29	0.00	96.4	0.63	0.10	0.51	0.69
6	0.41	0.18	0.48	0.20	0.51	1.01	97.1	0.00	0.31	0.00
7	0.20	0.09	1.07	0.59	0.31	0.22	0.10	95.7	0.62	1.98
8	0.31	0.18	1.45	0.99	0.81	0.78	0.73	0.29	95.6	1.88
9	0.00	0.00	0.10	0.10	0.81	0.11	0.00	0.49	0.21	92.0

Experiment

Prediction \ Label	0	1	2	3	4	5	6	7	8	9
0	97.4	0.00	0.58	0.00	0.10	0.22	0.84	0.10	0.00	0.40
1	0.00	98.0	0.68	0.00	0.10	0.00	0.21	0.39	0.10	0.50
2	0.10	0.18	92.3	0.30	0.31	0.00	0.42	2.04	0.51	0.20
3	0.10	0.26	1.26	95.1	0.00	1.12	0.10	0.49	0.41	0.59
4	0.00	0.00	0.58	0.10	96.7	0.11	0.31	0.29	0.41	1.98
5	1.02	0.26	0.10	1.19	0.00	95.0	0.52	0.19	1.03	1.19
6	0.82	0.35	0.87	0.10	0.51	0.78	96.7	0.00	0.21	0.10
7	0.10	0.09	0.68	0.89	0.20	0.00	0.10	95.2	0.51	2.28
8	0.51	0.79	2.81	2.28	1.32	2.69	0.84	0.68	96.7	2.97
9	0.00	0.00	0.10	0.10	0.71	0.11	0.00	0.58	0.10	89.8

Fig. S22. Confusion matrices showing the prediction results for the calculated and experimental convolutional neural networks.

References

1. Li, X. *et al.* Fast and reliable storage using a 5 bit, nonvolatile photonic memory cell. *Optica* **6**, 1 (2019).
2. Youngblood, N. *et al.* Tunable Volatility of Ge2Sb2Te5 in Integrated Photonics. *Adv. Funct. Mater.* **29**, (2019).
3. Rios, C. *et al.* Integrated all-photonic non-volatile multi-level memory. *Nat. Photonics* **9**, 725–732 (2015).
4. Rios, C. *et al.* Memcomputing on a photonic platform. *Submitted* (2017).
5. Li, X. *et al.* Fast and reliable storage using a 5 bit, nonvolatile photonic memory cell. *Optica* **6**, 1 (2019).
6. Feldmann, J. *et al.* Calculating with light using a chip-scale all-optical abacus. *Nat. Commun.* **8**, (2017).
7. Ma, Y. *et al.* Ultralow loss single layer submicron silicon waveguide crossing for SOI optical interconnect. *Opt. Express* **21**, 29374 (2013).
8. Lu, L., Wu, J., Wang, T. & Su, Y. Compact all-optical differential-equation solver based on silicon microring resonator. *Front. Optoelectron.* (2012) doi:10.1007/s12200-012-0186-9.
9. Zheng, J. *et al.* Nonvolatile electrically reconfigurable integrated photonic switch. *arXiv* 1–30 (2019).
10. Jouppi, N. P. *et al.* In-Datacenter Performance Analysis of a Tensor Processing Unit. *Proc. ISCA '17* (2017) doi:10.1145/3079856.3080212.
11. Sun, J., Timurdogan, E., Yaacobi, A., Hosseini, E. S. & Watts, M. R. Large-scale nanophotonic phased array. *Nature* **493**, 195–199 (2013).
12. Nahmias, M. A. *et al.* Photonic Multiply-Accumulate Operations for Neural Networks. *IEEE J. Sel. Top. Quantum Electron.* (2019) doi:10.1109/jstqe.2019.2941485.
13. Chen, L. & Lipson, M. Ultra-low capacitance and high speed germanium photodetectors on silicon. *Opt. Express* **17**, 7901 (2009).

14. Alphacore. ADC (Analog-to-Digital converters) – Alphacore.

15. Miller, D. Device Requirements for Optical Interconnects to Silicon Chips. *Proc. IEEE* **97**, 1166–1185 (2009).

16. Gehring, H., Eich, A., Schuck, C. & Pernice, W. H. P. Broadband out-of-plane coupling at visible wavelengths. *Opt. Lett.* **44**, 5089 (2019).

17. Oskooi, A. F. *et al.* Meep: A flexible free-software package for electromagnetic simulations by the FDTD method. *Comput. Phys. Commun.* **181**, 687–702 (2010).

18. Nvidia. Inference Platforms for HPC Data Centers | NVIDIA Deep Learning AI.

19. Merolla, P. A. *et al.* A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science (80-.).* **345**, 668–673 (2014).

20. Benjamin, B. V. *et al.* Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations. *Proc. IEEE* (2014) doi:10.1109/JPROC.2014.2313565.

21. Ambrogio, S. *et al.* Equivalent-accuracy accelerated neural-network training using analogue memory. *Nature* **558**, 60–67 (2018).

22. Li, C. *et al.* Analogue signal and image processing with large memristor crossbars. *Nat. Electron.* **1**, 52–59 (2018).

23. Shastri, B. J. *et al.* Principles of Neuromorphic Photonics. in *Encyclopedia of Complexity and Systems Science* (2018). doi:10.1007/978-3-642-27737-5_702-1.

24. Shen, Y. *et al.* Deep learning with coherent nanophotonic circuits. *Nat. Photonics* **11**, 441–446 (2017).

25. Peng, H. T. *et al.* Neuromorphic Photonic Integrated Circuits. *IEEE J. Sel. Top. Quantum Electron.* **24**, 1–16 (2018).