

<https://doi.org/10.1038/s42005-024-01549-1>

Fiber optic computing using distributed feedback

Check for updates

Brandon Redding¹ ✉, Joseph B. Murray¹, Joseph D. Hart¹, Zheyuan Zhu², Shuo S. Pang² & Raktim Sarma³

The widespread adoption of machine learning and other matrix intensive computing algorithms has renewed interest in analog optical computing, which has the potential to perform large-scale matrix multiplications with superior energy scaling and lower latency than digital electronics. However, most optical techniques rely on spatial multiplexing, requiring a large number of modulators and detectors, and are typically restricted to performing a single kernel convolution operation per layer. Here, we introduce a fiber-optic computing architecture based on temporal multiplexing and distributed feedback that performs multiple convolutions on the input data in a single layer. Using Rayleigh backscattering in standard single mode fiber, we show that this technique can efficiently apply a series of random nonlinear projections to the input data, facilitating a variety of computing tasks. The approach enables efficient energy scaling with orders of magnitude lower power consumption than GPUs, while maintaining low latency and high data-throughput.

Neural network-based machine learning algorithms have advanced dramatically in the last decade and are now routinely applied to a wide range of applications. However, this increase in performance has been accompanied by rapidly increasing computing demand—particularly in terms of the energy required to train and run these algorithms¹. This has inspired research in alternative platforms capable of performing the computationally-intensive matrix-vector multiplications (MVMs) and kernel convolution operations at the heart of most machine learning algorithms more efficiently. Among these approaches, optical computing is particularly promising due to its superior energy scaling and potential to overcome memory-access bottlenecks^{2–6}.

These unique features have led to a series of impressive demonstrations in which photonic computing systems have performed benchmark tasks with comparable accuracy to digital electronic neural networks while consuming orders of magnitude less energy^{3,7,8}. Most of these photonic computing schemes rely on spatial multiplexing in which input data is encoded in parallel on an array of modulators and the MVM output is recorded on an array of photodetectors. This approach has been explored both in free-space and integrated photonic platforms. Free-space platforms typically employ spatial light modulators (SLMs) to encode data and cameras to record the computed output. While this enables very large-scale computing (e.g. input vectors as large as $N \sim 10^6$ ⁹), the latency is limited by the SLM and camera speeds (in systems where the SLM is used to encode the input data). Recent demonstrations have shown that analog data from the scene can be

processed directly by a free-space, all-optical neural network^{10,11}, overcoming this latency constraint, although this approach is not easily applied to arbitrary and complex datasets (e.g. non-optical datasets such as the SONAR data analyzed in this work). Integrated photonic solutions are both more compact and have the potential for high-speed operation by exploiting state-of-the-art modulators and detectors^{12,13}. However, processing large matrices on-chip remains a challenge due to the size, heat, and complexity of integrating large numbers of individually addressable modulators. These limitations have inspired recent proposals for frequency^{14,15} and temporally multiplexed architectures in which a single modulator is used to encode an entire vector^{16–18}. These schemes could fill a gap in the photonic computing design space between the relatively slow, but large-scale free-space computing platforms that rely on SLMs to encode the input data and the high-speed but smaller-scale integrated photonic approaches. While initial demonstrations using temporal multiplexing were limited to a single neuron^{16,19}, spatio-temporal schemes were recently introduced that support larger scale networks^{8,18}, highlighting the potential for this approach to address the limitations of purely spatial multiplexing.

In addition, most of these photonic approaches are unable to natively perform multiple convolutions on the same input data. Convolutions have been used in a variety of machine learning techniques (e.g. convolutional neural networks) due to their ability to extract hierarchical features in a dataset^{20,21}. To perform multiple convolutions, photonic computing platforms need to generate multiple copies of the input data, separately apply a

¹US Naval Research Laboratory, Washington, DC, USA. ²CREOL, The College of Optics and Photonics, University of Central Florida, Orlando, FL, USA. ³Center for Integrated Nanotechnologies, Sandia National Laboratories, Albuquerque, NM, USA. ✉e-mail: brandon.redding@nrl.navy.mil

kernel transform to each copy, and record each result. For example, this has been achieved using free-space optics by applying a series of convolutions in the Fourier plane of an image using a diffractive optical element designed to map the result of different convolutions to separate positions on a camera²². Recently, a time-wavelength multiplexing scheme was proposed to address this challenge²³. This technique performed convolutions in parallel using optical frequency combs by encoding distinct kernel transforms on different sets of comb teeth. After wavelength de-multiplexing, an array of photodetectors was used to record the output of each kernel transform, enabling impressive throughput at the cost of increased system complexity and limited scalability (the number of kernels operations was limited by the number of individually addressable comb teeth).

In this work, we introduce a temporally multiplexed optical computing platform that performs convolutions using a simple and scalable approach based on distributed feedback in single mode fiber. We first encode the input vectors in the time domain as a pulse-train using a single high-speed modulator. This pulse-train is then injected into an optical fiber where a series of partial reflectors provide distributed feedback, generating a series of delayed copies of the input vector each weighted by the strength of a different reflector. In this demonstration, we rely on Rayleigh backscattering in standard single-mode fiber to provide this distributed feedback. Each Rayleigh scattering center creates a delayed copy of the input vector with random amplitude and phase—corresponding to the weights of a transformation matrix (i.e. a random kernel). The backscattered light is then recorded on a single, high-speed photodetector, performing the accumulation operation and introducing a non-linear transform. As explained in detail below, if the fiber is longer than the equivalent length of the encoded pulse train, then the fiber can perform multiple, distinct kernel operations on the input vector without requiring any additional routing or re-encoding of the input data. In principle, the weights of the transformation matrices (i.e. the kernels) could be inverse designed for specific computational tasks. As proof-of-concept, in this work, we use multiple random transformations to compute a non-linear random projection of the input vector into a higher dimensional space. We show that applying multiple random projections on the same input data can accelerate a variety of computing tasks including both unsupervised learning tasks such as non-linear principal component analysis (PCA) and supervised tasks including support vector machines (SVM) and extreme learning machines (ELM). More generally, this

approach offers 5 major benefits: (1) It natively performs multiple convolutions in parallel on the input data. (2) It is scalable and is capable of processing relatively large-scale matrices (we demonstrate matrix operations on vectors with 784 elements) while maintaining relatively high-speed (10 μs per MVM). (3) Since this approach relies on a passive transform to perform the matrix operations, the energy consumption scales as $O(N)$, enabling significant reductions in power consumption compared to a graphics processing unit (GPU). (4) The entire system is constructed using commercially available, fiber-coupled components, enabling a robust and compact platform. (5) Since the system operates directly on fiber-coupled, time-series data, this approach could be used to directly analyze data transmitted over fiber, opening up additional applications in remote sensing, RF photonics, and telecommunications.

Results

Operating principle

The basic operating principle is outlined in Fig. 1. Data is first encoded in the time domain as a series of optical pulses. This pulse train is then injected into an optical fiber where it is partially reflected by a series of Rayleigh scattering centers. This distributed backscattering process randomly mixes the elements in the input vector, resulting in a backscattered speckle pattern that contains a series of random projections of the input vector. The Rayleigh backscattered (RBS) light is then recorded on a high-speed photodetector which performs a non-linear transform on the backscattered electric field^{24,25}. In this work, we relied entirely on the non-linearity introduced by the detection process, although all-optical non-linearities introduced by light-matter interactions in optical fibers could be considered in the future. The digitized speckle pattern can then be used for a variety of computing tasks including non-linear principal component analysis, support vector machines, or extreme learning machines. More broadly, the Rayleigh backscattering process can form the first layer of an artificial neural network, efficiently expanding the input data to higher dimensional space. Since the specific weights and connections in the first layers of a neural network are not critical in most applications²⁶, this process can be used to accelerate one of the most computationally intensive tasks in a neural network. A digital electronics back-end can then be used to complete the neural network. In recent years, random projections have been proposed for a variety of computing tasks in both free-space^{9,24,25,27,28} and integrated photonic




Fig. 1 | Operating Principle. Data is encoded in the time-domain as a train of optical pulses (represented as blue dots in the equivalent neural network diagram). This pulse train is then injected into an optical fiber where distributed feedback is mediated by Rayleigh backscattering. The distributed feedback produces a series of delayed copies of the original data with random phase and amplitude (the yellow dots in the neural network diagram). Recording the backscattered signal on a

photodetector then provides a non-linear transform. Overall, this process results in a random non-linear projection of the original data into a higher dimensional space, facilitating a variety of computing and data analysis tasks including non-linear principal component analysis (PCA), support vector machine (SVM), or extreme learning machine (ELM).

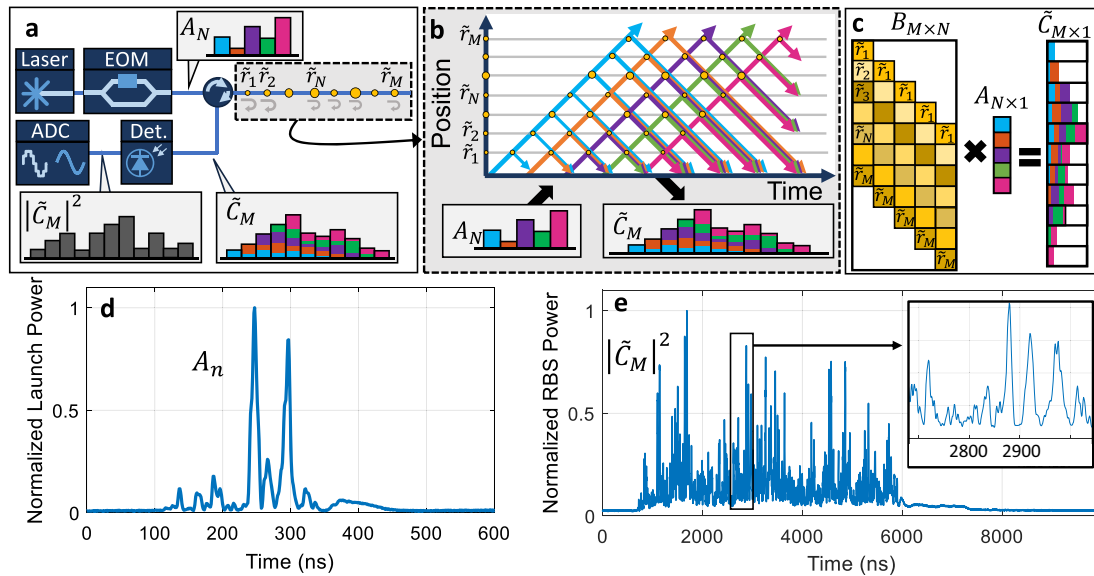


Fig. 2 | Experimental architecture. **a** The input vector A_N is encoded in the time domain using an EOM as a train of optical pulses which are injected into the fiber through a circulator. Rayleigh scattering (described by a series of complex reflectivity coefficients, \tilde{r}_i) provides distributed feedback and the backscattered field, \tilde{C}_M , is then recorded on a photodetector, providing a non-linear response. **b** The distributed scattering process can be described using a space-time diagram which tracks the position of each pulse as it travels through the fiber. The pulses (shown in different colors for clarity) are partially reflected by Rayleigh scattering centers with varying

complex reflectivity as they propagate down the fiber. As a result, the backscattered field contains randomly weighted contributions from each input pulse (i.e. each element in the input vector). **c** The distributed scattering process can be expressed as the multiplication between a complex transfer matrix B and the input vector A . **d** An example pulse train representing the SONAR data discussed in the SVM section and **e** the resulting RBS pattern. The inset in **e** shows a magnified view of the RBS pattern. EOM electro-optic modulator, SVM support vector machine, ADC analog-to-digital converter, Det photodetector, RBS Rayleigh backscattering.

platforms^{29,30}. While counterintuitive, researchers have shown that random transforms maintain key features in a dataset such as orthogonality while facilitating data analysis tasks such as dimensionality reduction or compressive sensing^{31,32}.

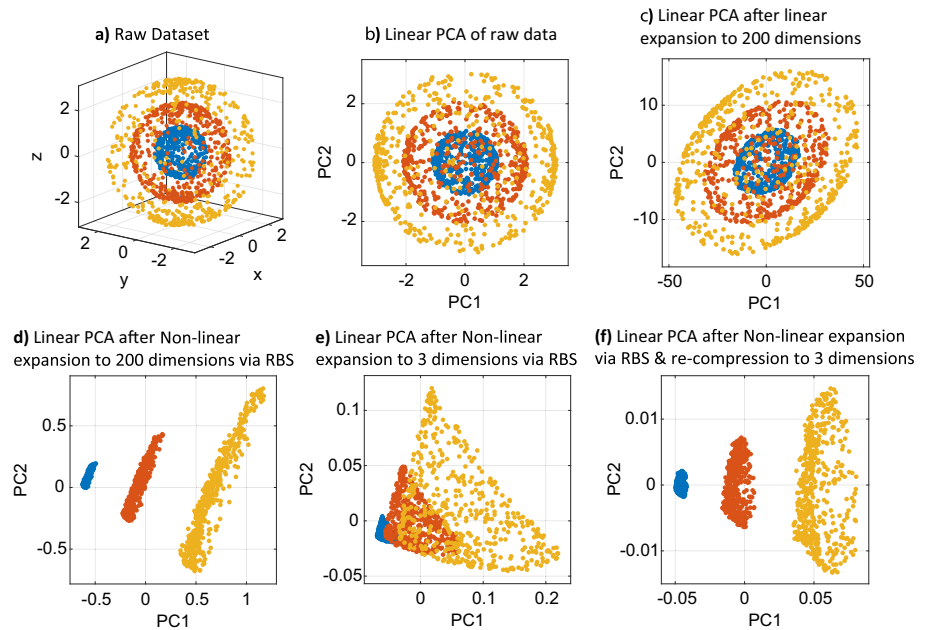
Experimentally, this basic approach was realized using the architecture shown in Fig. 2a. A continuous wave (CW) laser was coupled into an electro-optic modulator (EOM) which was used to encode the input data. An N -element vector, A_N , was encoded in the amplitude of a train of N pulses, as shown in the inset of Fig. 2a. This pulse train was coupled through a circulator into standard single-mode optical fiber where it was partially reflected by a series of Rayleigh scattering centers, creating a series of time-delayed copies of the original vector with random (though fixed) weights. The backscattered field, represented by an M -element complex vector \tilde{C}_M , was directed to a photodetector which performed a non-linear transform, generating photocurrent proportional to $|\tilde{C}_M|^2$. We used a narrow-linewidth (<2 kHz) seed laser to ensure that this coherent summation produced a repeatable speckle pattern. Figure 2d, e shows examples of an encoded pulse train and the resulting Rayleigh backscattered speckle pattern, illustrating the dramatic increase in dimensionality provided by the Rayleigh backscattering process. In this case, the input data was a 60-element vector ($N = 60$) representing a SONAR signal (discussed in the SVM section below) and the backscattered pattern was a ~2000-element vector.

The random kernel transform introduced by the Rayleigh backscattering process can be visualized using the space-time diagram shown in Fig. 2b. This diagram tracks the position of each pulse over time as it travels to the end of the fiber along a diagonal path in the upper-right direction. As each pulse propagates, it is partially reflected by a series of Rayleigh scattering centers and the reflected light travels back to the beginning of the fiber (along a diagonal path toward the bottom-right of the diagram). The Rayleigh backscattering at a given position in the fiber can be described by a complex reflectance, \tilde{r}_k , which is random but fixed. As shown in Fig. 2b, the backscattered light at time m includes contributions from each input pulse (once all of the pulses have entered the fiber) weighted by the complex reflectance at different positions in the fiber. While Fig. 2b presents a

simplified description of the Rayleigh backscattering process as a series of discrete partial reflectors, in reality, Rayleigh scattering is effectively continuous. However, the temporal correlation width of the Rayleigh backscattered light is set by the bandwidth (or pulse duration) of the encoded data³³. Thus, for the 5 ns pulses used in this work, the Rayleigh backscattering process can be approximated as a series of discrete complex reflectance coefficients spaced every 0.5 m in the fiber (the round-trip distance covered in 5 ns). Using this approximation, we can express the backscattered light at time m as the sum of each input vector element scaled by the appropriate reflection coefficient: $\tilde{C}_m = \sum_{n=1}^N A_n \tilde{r}_{m-n+1}$, where $\tilde{r}_k = 0$ for $k < 1$ or $k > M$. In other words, the Rayleigh backscattering process performs a series of vector convolution operations, applying different random kernels to the input vector as the pulse train propagates down the fiber. We can also express this transform as a matrix vector multiplication: $\tilde{C}_M = \tilde{B}_{M \times N} \times A_N$, where $\tilde{B}_{M \times N}$ is an $M \times N$ transfer matrix. As shown in Fig. 2c, the matrix B contains the reflection coefficients, \tilde{r}_k , arranged so that each row in \tilde{B} contains the same elements as the previous row, shifted by one column.

This system becomes particularly interesting if the fiber is longer than the equivalent length of the encoded pulse train (setting $M > N$). In this case the fiber can perform multiple, distinct random kernel operations on the same input vector. This is significant from a neural network perspective since convolution operations are known to help extract hierarchical features in a dataset^{20,21}. This is mathematically very different from most operations that have been implemented photonically in the past. First, photonic computing platforms exploiting random transforms^{9,29,34} performed global, fully-random projections of the input data, rather than convolutions. Second, most photonic computing platforms, including highly reconfigurable integrated photonic systems^{12,13}, are limited to applying a single kernel in each layer, rather than performing multiple, distinct kernel transforms on the same input data. A notable exception is the time-wavelength multiplexed approach which leveraged frequency combs to perform multiple convolutions in parallel²³. However, in addition to the complexity of this approach, the number of kernels was limited by the number of individually addressable comb teeth τ .

Fig. 3 | Non-linear principal component analysis (PCA) using distributed feedback. **a** The raw dataset consisted of 500 3-dimensional points randomly distributed on 3 concentric spheres. The colors (blue, red, orange) indicate data points on the 3 different spheres with radii of 1, 2, or 3. **b** Linear PCA of the raw data fails to separate the 3 classes of points. Plotted are the weights of the first 2 principle components (PC) for each datapoint. **c** A random linear transform is similarly unable to separate the 3 classes of points. **d** After the application of a non-linear random projection using Rayleigh backscattering, the 3 classes are clearly separable using a standard PCA. **e** Result of a PCA applied to 3 speckle grains selected from the RBS pattern, showing that the classes are difficult to separate without expanding the dimensionality. **f** Result of a PCA applied after re-compressing the RBS speckle pattern into 3 dimensions using a low-pass filter.



In the distributed feedback system, the number of kernel operations can be increased simply by using a longer optical fiber or a faster data encoding rate. In particular, the length of the output vector \tilde{C}_M is set by the length of the fiber, L , and the data encoding rate, f_0 , as $M \approx 2 \cdot \lceil (2L/(c/n))/\tau + N \rceil$, where c is the speed of light, n is the effective index in the fiber, and $\tau = 1/f_0$ is the length of the pulses representing the elements of A_N . The term in the square brackets represents the round-trip time in the fiber while the factor of 2 outside the brackets accounts for using a polarization diversity receiver to record the backscattered light in both polarizations in parallel (see Methods). The additional N accounts for the length of the input vector and assumes we make use of RBS light that does not include every element in A_N (i.e. RBS light collected before the entire pulse train enters the fiber and after the pulse train starts to leave the fiber). This expression also assumes that backscattered light is sampled at the data encoding rate of f_0 and that the temporal correlation width of the Rayleigh pattern matches the pulse duration τ , which is the case for Rayleigh scattering³³. By increasing the data encoding rate and the fiber length, this technique could be used to process large scale matrices or perform multiple distinct kernel operations on the same input vector (the number of distinct kernels is set by the ratio M/N). For example, this platform could perform MVMs with $M = 10^6$ using an encoding rate of 10 GHz and a 5 km fiber before attenuation becomes significant (note that 0.2 dB/km is typical for telecom fiber at a wavelength of 1550 nm, resulting in a round-trip loss of 2 dB for a 5 km fiber).

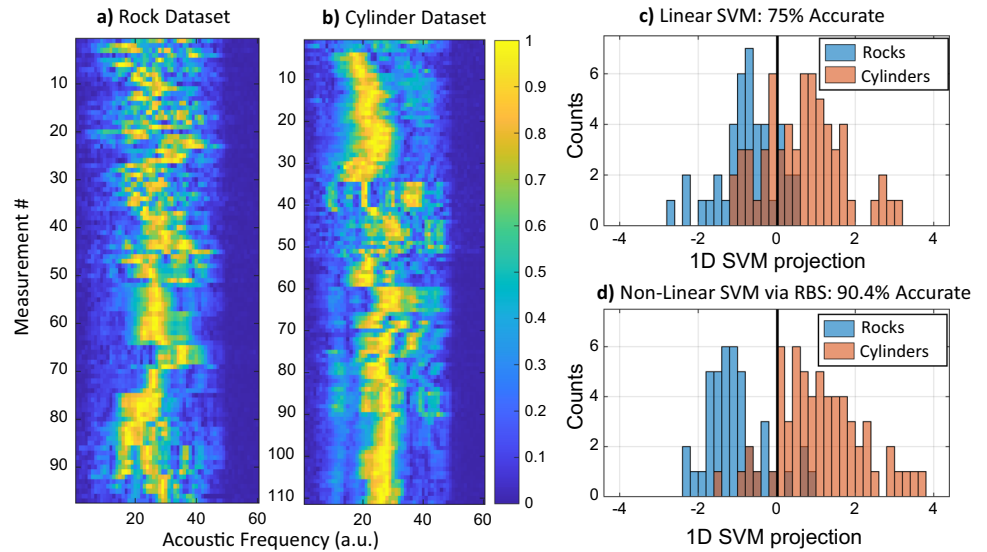
Although the temporal multiplexing approach presented here trades-off computing speed for the ability to use a single modulator and detector, the availability of high-speed optical modulators and detectors (e.g. 20 GHz devices are widely available) helps to mitigate this trade-off. The time required to compute a MVM can be expressed as $\tau_{\text{MVM}} = N/f_0 + 2L/(c/n)$, where the first term accounts for the length of the input pulse train and the second term represents the round-trip time in the fiber. In this work, we used a 500 m fiber and a 200 MHz encoding rate, yielding $\tau_{\text{MVM}} \approx 10 \mu\text{s}$ for N approaching 10^3 . Increasing the encoding rate to 20 GHz and using a 5 m fiber could enable a 100× speed-up ($\tau_{\text{MVM}} \approx 100 \text{ ns}$) while performing a MVM with the same matrix dimensions. This analysis implies that using higher frequency encoding is generally beneficial, enabling faster computation for a given matrix size. However, as we will discuss below, the power consumption also increases with the encoding rate and the proper balance will depend on the application.

Non-linear principal component analysis

As an initial demonstration of this technique, we considered a text-book example of non-linear principal component analysis. Non-linear PCA is an unsupervised learning technique that has been used for dimensionality reduction, singular value decomposition, denoising, and regression analysis^{35,36}. A standard PCA relies on linear transforms to project data onto a new coordinate system that represents the variance in a dataset using as few dimensions as possible. However, relying entirely on linear transforms limits a standard PCA to analyzing data that is linearly separable³⁷. A non-linear PCA operates by first applying a non-linear transform to a dataset before performing a standard PCA, facilitating the analysis of a wider range of data types. To illustrate how our platform can be used for non-linear PCA, we first created a dataset consisting of 500 points (defined by their cartesian coordinates x, y, z) randomly distributed on 3 concentric spheres with radii of 1, 2, or 3, as shown in Fig. 3a. Ideally, the PCA would decompose the output vector into a single non-zero principal component representing the length of the x, y, z vector. At minimum, the PCA should result in <3 significant PCs representing a low dimensional space in which the data points can be linearly separated. A standard linear PCA is unable to separate these 3 classes of points, as shown in Fig. 3b which plots the weights of the first 2 PCs for each data point. Moreover, simply expanding the dimensionality without applying a non-linear transform was unable to separate the classes on its own. To illustrate this, we computationally applied a linear random transform to each coordinate by multiplying each coordinate by a 3×200 random matrix before applying a PCA. As shown in Fig. 3c, this cascaded transform (linear dimension expansion followed by PCA) is still linear and is unable to separate the classes.

To perform a non-linear PCA, we used the Rayleigh backscattering platform to create a non-linear projection of the data onto a higher-dimensional space before applying a PCA. To do this, we injected each data point into the fiber using a train of three 50 ns pulses (i.e. an input vector with $N = 3$ and $\tau = 50 \text{ ns}$). The EOM was initially biased at zero transmission and the voltage sent to the EOM was set by the amplitude of the x, y, z coordinates of each point (normalized such that the maximum coordinate of 3 was set to the maximum transmission voltage for the EOM, V_π). We then recorded the Rayleigh backscattered speckle pattern produced by each point, yielding a $M = 200$ output vector. This process effectively projected each 3-dimensional point into a 200-dimensional speckle pattern.

Fig. 4 | Non-linear Support Vector Machine (SVM) using distributed feedback. The SONAR dataset consisted of 97 measurements of rocks (a) and 111 measurements of cylinders (b). c Applying a linear SVM directly to the training data resulted in a classification accuracy of 75%. d After using the optical platform to perform a non-linear random projection on the SONAR data, the SVM accuracy increased to 90.4%.



We then performed a standard linear PCA on the backscattered speckle patterns. As shown in Fig. 3d, the hybrid photonic/electronic non-linear PCA efficiently separated the three classes of points. This separation relied on both expanding the dimensionality and applying a non-linear transform to the original data. To illustrate this, we also attempted to perform a PCA using just 3 of the speckle grains in each backscattered speckle pattern. As shown in Fig. 3e, this was unable to efficiently separate the three classes (the separation depends on the 3 speckle grains selected, but the result shown is typical for most sets of 3 speckle grains). Expanding the dimensionality increased the chance of randomly finding a transform that projected the data into a space where it is highly separable.

While this demonstration showed that non-linear random projections can help identify variations in a dataset, the benefit of using analog optics (in terms of power consumption and speed) was limited since the final PCA was performed on a relatively large dataset with 200 dimensions. Fortunately, after expanding the dataset with the non-linear transform, we could then compress it into a lower dimensional vector in the analog optical domain before performing the final PCA while maintaining many of the benefits. To test this, we re-compressed the 200-dimensional RBS speckle pattern produced by each point to 3-dimensions by averaging 100 points at a time. Experimentally, this could be achieved using a photodetector with a low-pass filtering response. As shown in Fig. 3f, a PCA performed on the re-compressed vectors was still able to efficiently separate the three classes. This simple demonstration illustrates how a non-linear random projection can be used to identify key features in a dataset.

Non-linear support vector machine

The same non-linear random projections can be used to construct a non-linear support vector machine (NL-SVM). SVMs are a supervised learning technique designed to separate different classes in a dataset by identifying the maximum margin hyperplane separating two classes. A non-linear SVM first projects the data into a higher-dimensional space before finding a hyperplane to separate different classes^{38,39}. The platform proposed in this work is ideally suited for non-linear SVM, since the convolutions introduced by the RBS process can efficiently project an input dataset into a higher-dimensional space to facilitate classification.

To explore the use of our platform for NL-SVM, we selected a benchmark SONAR dataset consisting of SONAR measurements of either rocks or metal cylinders (phantoms for underwater mines)⁴⁰. The dataset consists of 97 measurements of rocks and 111 measurements of cylinders. Each measurement was obtained using a frequency-modulated

chirped SONAR and contains a 60-element vector representing the reflected signal as a function of acoustic frequency, as shown in Fig. 4a, b. This dataset has been divided into a training and a validation dataset, each containing 104 measurements (we focused on the Aspect Angle Dependent Test described in⁴⁰). We first attempted to classify this data by using a standard linear SVM to assign a hyperplane based on the training data and then evaluated how well this hyperplane separates the validation data. A histogram showing the position of the validation data measurements along a 1-dimensional SVM projection is shown in Fig. 4c. Ideally, the two classes would be completely separable in this space; however, we observed significant overlap between the classes and obtained an overall classification accuracy of only 75%. We then used the RBS platform to perform a random non-linear projection on the SONAR data before applying an SVM. In this case, each SONAR measurement (consisting of a 60-element vector) was encoded in a pulse train using 5 ns pulses. The RBS process was used to expand the dimensionality of the SONAR signal from 60 to 2000. An example of an encoded pulse train and the resulting RBS pattern are shown in Fig. 2d, e. As shown in Fig. 4d, the non-linear SVM was much more effective at separating the two classes, obtaining an accuracy of 90.4% on the validation data, comparable to the performance of the neural network reported in⁴⁰. This illustrates the potential for our platform to facilitate data classification by transforming input data into higher dimensional space using random convolutions.

Extreme learning machine

Extreme learning machines are a type of feed-forward neural network in which the weights and connections in the hidden layers are fixed and a single decision layer is trained to complete a task^{41,42}. ELMs were initially proposed to avoid the computational demands of training every connection in a neural network, but their unique structure is particularly well-suited for optical implementations. Photonic ELMs can use passive photonic structures to apply a complex transform (i.e. the fixed layer in the ELM) and rely on a single electronic decision layer to complete the computing task^{28,43}. This has enabled photonic computing architectures built around multimode fiber²⁸ or complex disordered materials²⁴ where precise control of the transfer matrix would be challenging. Here, we show that our distributed optical feedback platform can be configured as an ELM to perform image classification.

We tested our system on two benchmark tasks: classifying the MNIST Digit database and the MNIST Fashion database⁴⁴. Each dataset consists of 60,000 training images and 10,000 test images in 10 classes (either hand-written digits from 0-9 or 10 types of clothing). Although

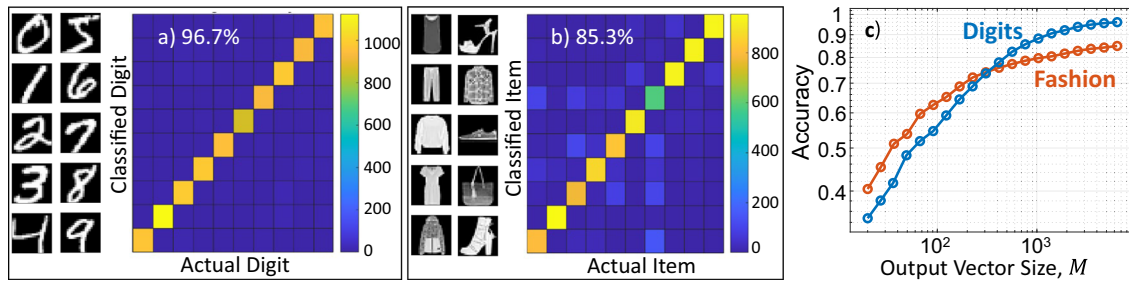


Fig. 5 | Extreme learning machine (ELM) using distributed feedback. **a, b** The photonic ELM classified the MNIST digit database with an accuracy of 96.7% and the Fashion database with an accuracy of 85.3%. **c** The ELM accuracy increases with the dimensionality of the output vector, M .

the distributed feedback system performs vector convolutions, it can also be used to perform two-dimensional convolutions on image data that has been flattened to a one-dimensional vector (see Methods for details)²³. To do this, we encoded each image as a one-dimensional pulse train by assigning the magnitude of each 5 ns pulse to the intensity of one pixel in the image. We injected pulse trains representing all 70,000 images in series and recorded 70,000 Rayleigh backscattered speckle patterns. We then used a ridge regression algorithm to train a decision layer to classify the 60,000 training images. Finally, we tested the ELM using the backscattered patterns obtained from the 10,000 validation images. This same process was repeated for the digit and fashion databases.

As shown in Fig. 5a, b, the distributed feedback ELM system achieved an accuracy of 96.7% and 85.3% on the digit and fashion databases, respectively. This performance is comparable to the best performing photonic neural networks. For example, on the digit database, a free-space optical extreme learning machine achieved an accuracy of 92%⁴³, similar to a diffractive optical neural network which achieved 91.75%¹¹. The time-wavelength multiplexed convolutional accelerator reported an accuracy of ~88%²³. Integrated photonic approaches also report similar accuracies, including 95.3%¹³ for a wavelength multiplexed scheme and 93.1%⁸ using a VCSEL array. Hybrid optical/electronic approaches have yielded accuracy as high as 98%⁴⁵. Our results are also similar to digital electronic neural networks with similar depth (~96%), while multi-layer neural networks achieve an accuracy above 99%⁴⁴.

Our approach is able to achieve this accuracy simply by expanding the dimensionality and leveraging the inherent power of convolutions to extract high-level features from a dataset, without explicitly training or controlling the underlying transforms. This demonstration also shows that the same platform, and, in fact, the same optical fiber, can be used for multiple tasks simply by re-training the decision layer, consistent with previous demonstrations of the versatility of photonic ELMs^{28,43}.

The most important parameter impacting the ELM accuracy is the length of the Rayleigh backscattering pattern (i.e. the length of the output vector, or the extent to which the system expanded the dimensionality of the original data). The entire measured RBS pattern consisted of ~2500 speckle grains (a ~6× increase in dimensionality compared to the 400-pixel MNIST Digit images). To investigate this dependence, we subsampled the measured RBS pattern and repeated the training and test procedure using output vectors of varying length. As shown in Fig. 5(c), the accuracy increases rapidly with the length of the output vector before gradually plateauing. Higher accuracy is achieved for the digit database, which is consistent with previous studies showing that the fashion database is more challenging⁴⁴. Nonetheless, the accuracy increases monotonically with the size of the output dimension in both cases. Fortunately, the distributed feedback platform is well suited for this task and can increase the dimensionality of the output data simply by using a fiber that is longer than the effective length of the pulse train. Moreover, since the backscattering process is entirely passive, increasing the dimensionality in this way has a negligible effect on the power consumption despite increasing the size of the MVM.

Energy consumption

As discussed, the distributed feedback process effectively performs a series of random convolutions on the input vector, which can be expressed as a matrix vector multiplication. This framework allows us to compare the power consumption that could be achieved by our approach to a standard digital electronic processor like a GPU by analyzing the energy per multiply-accumulate (MAC) operation. Our analysis accounts for the energy required to operate the laser, the modulator, and the photodetector as well as the digital to analog converter (DAC) and analog to digital converter (ADC). Assuming shot-noise limited detection, we first estimated the optical power required at the detector to obtain a Rayleigh backscattering pattern with the desired signal-to-noise ratio (SNR), as⁴⁶:

$$P_{Rx} = 2^{2b} q f_0 / \mathcal{R} \tag{1}$$

where b is the effective number of bits (ENOB) with $\text{SNR} = 6.02 \times b + 1.76$ in dB, q is the charge of an electron, f_0 is the measurement bandwidth (and the data encoding rate), and \mathcal{R} is the responsivity of the detector. We can then estimate the required laser power as

$$P_{laser} = P_{Rx} / [T_{mod} r_{RBS}] \tag{2}$$

where T_{mod} is the transmission through the modulator and r_{RBS} is the average reflectance due to Rayleigh backscattering in the fiber. Since Rayleigh backscattering is a distributed process, the effective Rayleigh backscattering coefficient depends on the duration of the pulse train launched into the fiber and can be estimated as

$$r_{RBS} = (-82 \text{ dB/ns})(N \cdot \tau) \tag{3}$$

for Corning SMF-28e+ (as used in this work)⁴⁷, where N is the length of the input vector and τ is the pulse duration for each element in the input vector. The total electrical power required to operate the laser can then be calculated as P_{laser} / η , where η is the wall-plug efficiency of the laser.

The electrical power consumed by the modulator can be expressed as⁴⁸


$$P_{mod} = \frac{N}{M} \left(\frac{1}{2} C_{mod} V_{\pi}^2 f_0 \right) \tag{4}$$

where N is the length of the input vector, M is the length of the output vector (i.e. the Rayleigh backscattering pattern), C_{mod} is the capacitance of the modulator, and V_{π} is the peak-to-peak driving voltage of the modulator. The ratio of N/M denotes the operating duty-cycle of the modulator and we assume $M \geq N$, which corresponds to expanding the dimensionality of the input vector. The electrical power consumed by the photodetector can be expressed as

$$P_{det} \approx V_{bias} \mathcal{R} P_{Rx} \tag{5}$$

where V_{bias} is the bias voltage applied to the detector. State-of-the-art DACs and ADCs require ~0.5 pJ/use and ~1 pJ/use, respectively⁴⁹⁻⁵¹. The DAC is

Fig. 6 | Power consumption. **a** The energy/MAC required by the laser, detector, modulator, DAC, and ADC as a function of input vector size at an encoding rate of 200 MHz and assuming an ENOB of 6. **b** The total energy/MAC assuming an ENOB of 6 for varying encoding rates, f_0 , using either standard fiber (i.e. relying on Rayleigh backscattering, RBS) or using enhanced Rayleigh backscattering fiber with 15 dB higher reflectance. GPU graphics processing unit, ADC analog-to-digital convert, DAC digital-to-analog converter, RBS Rayleigh backscattering, MAC multiply-accumulate, ENOB effective number of bits.



used each time we switch the modulator, which occurs $(N/M)f_0$ times per second, whereas the ADC operates continuously at a rate of f_0 , resulting in power consumption of

$$P_{\text{DAC}} = 0.5 \text{ pJ} \times \left(\frac{N}{M}\right) f_0 \quad (6)$$

$$P_{\text{ADC}} = 1 \text{ pJ} \times f_0 \quad (7)$$

The total power can then be expressed as $P_{\text{total}} = P_{\text{laser}}/\eta + P_{\text{mod}} + P_{\text{det}} + P_{\text{DAC}} + P_{\text{ADC}}$. In order to estimate the energy per MAC, we can then calculate the total number of MACs per second as $N \cdot M/\tau_{\text{MVM}}$, where τ_{MVM} is the time required to complete a matrix vector multiplication. Since a MVM is completed in the time required to record a Rayleigh backscattering pattern with M outputs, $\tau_{\text{MVM}} = M \cdot \tau$. Thus, the number of MACs per second is N/τ or Nf_0 , and the energy per MAC can be estimated as $P_{\text{total}}/(Nf_0)$. Based on this analysis, we find that the energy per MAC required to power the laser scales as f_0/N^2 whereas the energy per MAC required to power the modulator, detector, DAC, and ADC scales as $1/N$ and is independent of modulation frequency. This difference in scaling results from the dependence of r_{RBS} on N and τ in Eq. 3. Finally, this analysis ignored the use of a polarization diversity receiver for simplicity. However, using polarization diversity does not change the energy per MAC. For example, using polarization diversity would require twice the laser power to achieve the same SNR, since the RBS light is split between two detectors, but enables twice the MACs per second. Instead, using a polarization diversity receiver doubles the data throughput (for the same modulation and detection speeds) without affecting the power efficiency.

This model allowed us to quantitatively evaluate the energy/MAC which could be achieved using the optical distributed feedback scheme. Here, we assumed a required ENOB of $b = 6$, which corresponds to a SNR of 38 dB, considerably higher than the 23.6 dB SNR of the RBS patterns measured experimentally in this work (see Methods). We set $M = 10 \times N$, implying that the system was used to expand the dimensionality of the input vector by a factor of 10. We then assumed typical values for the remaining parameters: a laser with wall-plug efficiency of $\eta = 0.2$, a modulator with $C_{\text{mod}} = 1\text{fF}$ and $V_{\pi} = 1\text{V}$, and a detector with $V_{\text{bias}} = 3\text{V}$ and $\mathcal{R} = 1\text{A/W}^{52,53}$. We first estimated the energy per MAC as a function of the input vector size N using a modulation frequency of 200 MHz, matching our experimental conditions. As shown in Fig. 6a, the energy/MAC is dominated by the laser under these conditions. For comparison, Fig. 6 also shows the power consumption typical of state-of-the-art GPUs (~ 1 pJ/MAC), which is independent of vector size³⁰. At a vector size of $N = 784$, corresponding to the size of the 28×28 pixel images in the fashion database, the RBS system provides a $30 \times$ reduction in power consumption compared to a GPU. This improvement increases dramatically for larger vector sizes due to the $1/N^2$ dependency of the laser energy. Reducing the required ENOB would also significantly reduce the power consumption by reducing the required laser power.

For many applications, using higher modulation frequencies could be attractive to enable low latency computing and higher data throughput. For example, increasing the modulation rate to 10 GHz would reduce the latency from $10 \mu\text{s}$ to 200 ns for the same size MVMs presented in the ELM Section. In terms of computing power, operating at 10 GHz would enable this system to perform $>10^{12}$ MACs/second assuming an input vector size of $N = 10^3$, approaching the throughput of time-wavelength multiplexed schemes which have achieved 11×10^{12} operations per second²³. The expected energy/MAC at an encoding rate of 10 GHz is shown in orange in Fig. 6(b) (the laser and ADC are still the dominant sources of energy consumption). In this system, the energy/MAC increases with encoding rate to compensate for the reduced backscattering at higher encoding rates (since r_{RBS} depends on τ). Nonetheless, the RBS system still outperforms a GPU for input vectors with $N > 10^3$ at an encoding rate of 10 GHz. In order to reduce the energy consumption, we could use enhanced backscattering fiber, which is commercially available and provides ~ 15 dB higher backscattering than standard single-mode fiber⁵⁴. Figure 6(b) also shows the energy/MAC required using enhanced backscattering fiber with a data encoding rate of 10 GHz, which approaches the energy/MAC achieved at 200 MHz using standard fiber. Combining enhanced scattering fiber with a 200 MHz encoding rate could enable even lower power consumption, outperforming a GPU for vectors with $N > 25$. Although this type of fiber was initially developed for fiber optic sensing, it is ideally suited for this type of optical computing platform and could enable a significant reduction in power consumption. A distributed feedback fiber employing ultra-weak fiber Bragg gratings⁵⁵ or point reflectors⁵⁶ could also be used to achieve higher reflectance and thus lower power consumption. These technologies could also enable customized kernel transforms by tailoring the position and strength of each reflector.

Discussion

One challenge not addressed in this initial work is the environmental stability of the fiber. The PCA, SVM, and ELM algorithms investigated here each assume that a given launch pattern will always produce the same Rayleigh backscattering pattern (albeit with a finite SNR). However, the Rayleigh backscattering pattern depends on temperature and strain in addition to the launch pattern (Rayleigh backscattering in optical fiber is often used for distributed temperature and strain sensing⁵⁷). As an example, for the $N = 400$ pixel images in the digit database encoded at 200 MHz, the Rayleigh backscattering pattern would decorrelate if the temperature of the fiber drifted by $0.4 \text{ m}^\circ\text{C}$ (see Methods for details). While this level of temperature stabilization is possible, single-mode fiber is a very controlled platform with limited degrees of freedom and relatively simple calibration techniques have been proposed in other applications leveraging Rayleigh backscattering which might be applicable here⁵⁸. Note that in the experiments reported in this work, the fiber was enclosed in an aluminum box, but was not temperature stabilized. Instead, the training and validation datasets were recorded during a short

measurement window while the fiber environment remained stable. For example, the ELM experiments required 0.7 s to record the RBS pattern from all 70,000 test images and we found that the RBS pattern did not drift significantly on this time-scale. Similarly, laser frequency drift could also impact the repeatability of the Rayleigh backscattering pattern. The sensitivity to laser frequency depends on the inverse of the pulse-train duration. For example, the digit database images were encoded using 2 μ s pulse trains, requiring frequency stability better than \sim 500 kHz. In this work, we used a narrow-linewidth (\sim kHz) laser which was sufficiently stable on the time-scale of these measurements. However, active frequency stabilization may be required for longer term operation. Alternately, faster encoding or smaller datasets would relax the frequency stability requirement.

While this work introduced a basic architecture for performing optical computing with distributed feedback, there are several areas for improvement and further investigation. First, in this work, we relied entirely on the photodetection process to introduce the required non-linearity. Introducing non-linearities induced by light-matter interaction in the optical domain could enable more sophisticated, multi-layer networks capable of addressing more challenging computing tasks⁵⁹. There are a few promising approaches which could be pursued to introduce intra-network non-linearities in this platform. For example, an analog optical-electrical-optical conversion could be introduced, following the approach used in single-node fiber-optic reservoir computers¹⁹ and recently explored in the context of integrated photonic neural networks^{59,60}. Alternately, the long interaction lengths and high-power density in single mode fiber could be leveraged to introduce all-optical non-linearities such as four-wave mixing, modulation instability, or stimulated Brillouin scattering⁶¹. Second, the temporal multiplexing scheme proposed here sacrifices computing speed for simplicity by enabling the use of a single modulator to encode the entire input vector. In the future, some degree of spatial multiplexing could also be included to balance this trade-off. On the detection side, this is particularly straightforward and was already implemented in part using a polarization diversity receiver. In the future, recording the backscattered pattern from a few-mode fiber or multiple discrete fibers could increase the output vector size without compromising the computing speed. Some degree of parallelization could also be explored on the data encoding side, e.g., by using multiple modulators and coupling light into a few-mode fiber. Third, while we focused on random projections in this work, the same basic combination of temporal encoding and distributed feedback could be used to perform other operations. For example, instead of relying on Rayleigh backscattering, a series of carefully positioned partial reflectors (such as weak fiber Bragg gratings⁵⁵ or point reflectors⁵⁶) could be used to implement distinct kernel transforms at different positions along the fiber (the reflectivity at each position in the fiber could then be set to encode a desired kernel weight). Alternately, inverse design principles could be used to optimize the reflector geometry for a desired kernel response. These reflectors could potentially be reconfigurable (e.g. by straining a weak fiber Bragg grating to adjust its reflectivity at the seed laser wavelength), enabling reconfigurable kernels.

In summary, this work introduced an optical computing platform based on temporal multiplexing and distributed feedback that performs random convolutions using a passive optical fiber. We showed that Rayleigh backscattering in single mode fiber can be used to perform non-linear random kernel transforms on arbitrary input data to facilitate a variety of computing tasks, including non-linear principal component analysis, support vector machines, or extreme learning machines. This approach enables large scale MVMs with $O(N)$ energy scaling using a single modulator and photodetector. The entire system can be constructed using off-the-shelf fiber-coupled components, providing a compact and accessible approach to analog optical computing. Finally, since this approach operates on temporally encoded light in standard single-mode fiber, it could potentially be applied directly to optical data transmitted over fiber enabling applications in remote sensing and RF photonics.

Methods

Experimental details

A narrow-linewidth laser (<2 kHz, RIO Orion, Grade 4) was used as the seed laser. In the experiments reported in this work, an Erbium-doped fiber amplifier (EDFA) was inserted after the EOM in Fig. 2 to increase the launch power and compensate for the use of a relatively low power (10 mW) seed laser and relatively high insertion loss EOM (4.5 dB). In the future, a higher power laser could be used to avoid needing the EDFA (depending on the length of the pulse train and encoding rate, peak launch power in the range of 10-100 mW is sufficient to achieve an *enob* of 6 based on Eqs. 1–3). In this work, we also used a second EDFA followed by a wavelength division multiplexing filter after the circulator to amplify the Rayleigh backscattered pattern before detection. This EDFA was used to minimize the effect of photodetector noise. In the future, using a lower-noise detector would preclude the need for the second EDFA. Note that the EDFA operated in the small signal gain regime and did not introduce a non-linear response to the RBS pattern. A polarization diversity receiver (a polarizing beam splitter combined with two detectors) was used to record the RBS pattern in orthogonal polarizations.

In the ELM experiments, the 2-dimensional images were flattened into 1-dimensional vectors. As discussed in²², vector convolutions can be used to perform a convolution on 2-dimensional image data, although there is an overhead cost (not all output measurements are used) and the stride is inherently asymmetric (a symmetric stride could be obtained through sequential measurements by modifying the encoding strategy). In the image recognition tasks described here, we used all of the output samples, which included measurements representing standard 2D convolutions as well as mixtures that convolved different combinations of input pixels and would not be computed in a standard convolution. We also did not attempt to correct for the asymmetric stride introduced by the vector convolution process.

The digitized RBS speckle patterns were recorded at 1 GS/s. This is slightly faster than the temporal correlation width, which is set by the data encoding rate (i.e. 200 MHz for the SVM or ELM test and 20 MHz for the PCA test)³³. In this work, we processed slightly oversampled Rayleigh backscattering patterns, but the dimensionality listed as M in each section corresponded to the pattern length divided by the correlation width of 5 ns or 50 ns.

We calculated the SNR of the RBS patterns recorded in this work as the ratio of the average power of the RBS pattern to the standard deviation in the background (recorded with the EOM blocked). We found an SNR of 23.6 dB, which corresponds to an *enob* of 3.6. Fortunately, neural networks are quite robust to low-precision computing⁶², enabling the excellent performance achieved in the benchmark tasks explored in this work despite this modest precision.

Temperature Sensitivity

The temperature sensitivity of the Rayleigh backscattering pattern depends on the length of the incident pulse train and can be estimated as $(f_0/N)[1.2 \text{ GHz}^\circ\text{C}]$, where the (f_0/N) term represents the transform limit of the input pulse train (which has a duration of $N\tau = N/f_0$) and the term in square brackets represents the shift in the Rayleigh spectrum with temperature⁶³. This expression calculates the temperature shift required for the Rayleigh backscattering spectrum to shift by the transform limit of the pulse train, which would result in a decorrelated speckle pattern.

Data availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Received: 19 September 2023; Accepted: 5 February 2024;

Published online: 02 March 2024

References

- Patterson, D. et al. Carbon emissions and large neural network training. Preprint at *arXiv* <https://doi.org/10.48550/arXiv.2104.10350> (2021).
- Zhou, H. et al. Photonic matrix multiplication lights up photonic accelerator and beyond. *Light Sci. Appl.* **11**, 30 (2022).
- Wetzstein, G. et al. Inference in artificial intelligence with deep optics and photonics. *Nature* **588**, 39–47 (2020).
- Solli, D. R. & Jalali, B. Analog optical computing. *Nat. Photonics* **9**, 704–706 (2015).
- Wu, J. et al. Analog optical computing for artificial intelligence. *Engineering* **10**, 133–145 (2022).
- Shastri, B. J. et al. Photonics for artificial intelligence and neuromorphic computing. *Nat. Photonics* **15**, 102–114 (2021).
- El Srouji, L. et al. Photonic and optoelectronic neuromorphic computing. *APL Photonics* **7**, 051101 (2022).
- Chen, Z. et al. Deep learning with coherent VCSEL neural networks. *Nat. Photonics* **17**, 723–730 (2023).
- Rafayelyan, M., Dong, J., Tan, Y., Krzakala, F. & Gigan, S. Large-scale optical reservoir computing for spatiotemporal chaotic systems prediction. *Phys. Rev. X* **10**, 41037 (2020).
- Wang, T. et al. Image sensing with multilayer nonlinear optical neural networks. *Nat. Photonics* **17**, 408 (2023).
- Lin, X. et al. All-optical machine learning using diffractive deep neural networks. *Science* **361**, 1004 (2018).
- Shen, Y. et al. Deep learning with coherent nanophotonic circuits. *Nat. Photonics* **11**, 441–446 (2017).
- Feldmann, J. et al. Parallel convolutional processing using an integrated photonic tensor core. *Nature* **589**, 52–58 (2021).
- Fan, L. et al. Multidimensional convolution operation with synthetic frequency dimensions in photonics. *Phys. Rev. Appl.* **18**, 034088 (2022).
- Fan, L., Wang, K., Wang, H., Dutt, A. & Fan, S. Experimental realization of convolution processing in photonic synthetic frequency dimensions. *Sci. Adv.* **9**, 4956 (2023).
- Appelant, L. et al. Information processing using a single dynamical node as complex system. *Nat. Commun.* **2**, 468 (2011).
- Stelzer, F., Röhm, A., Vicente, R., Fischer, I. & Yanchuk, S. Deep neural networks using a single neuron: folded-in-time architecture using feedback-modulated delay loops. *Nat. Commun.* **12**, 5164 (2021).
- 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).
- Larger, L. et al. High-speed photonic reservoir computing using a time-delay-based architecture: million words per second classification. *Phys. Rev. X* **7**, 011015 (2017).
- Lawrence, S., Giles, C. L., Tsoi, A. C. & Back, A. D. Face recognition: a convolutional neural-network approach. *IEEE Trans. Neural Networks* **8**, 98–113 (1997).
- Lecun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
- Chang, J., Sitzmann, V., Dun, X., Heidrich, W. & Wetzstein, G. Hybrid optical-electronic convolutional neural networks with optimized diffractive optics for image classification. *Sci. Rep.* **8**, 12324 (2018).
- Xu, X. et al. 11 TOPS photonic convolutional accelerator for optical neural networks. *Nature* **589**, 44–51 (2021).
- Ando, T., Horisaki, R. & Tanida, J. Speckle-learning-based object recognition through scattering media. *Opt. Express* **23**, 33902–33910 (2015).
- Saade, A. et al. Random projections through multiple optical scattering: approximating kernels at the speed of light. in *2016 IEEE Int. Conf. Acoust. Speech Signal Process.* 6215 (2016).
- Havasi, M., Peharz, R. & Hernandez-Lobato, J. M. Minimal random code learning: getting bits back from compressed model parameters. Preprint at *arXiv* <https://doi.org/10.48550/arXiv.1810.00440> (2018).
- Dong, J., Rafayelyan, M., Krzakala, F. & Gigan, S. Optical reservoir computing using multiple light scattering for chaotic systems prediction. *IEEE J. Sel. Top. Quantum Electron.* **26**, 7701012 (2020).
- Teğin, U., Yıldırım, M., Oğuz, İ., Moser, C. & Psaltis, D. Scalable optical learning operator. *Nat. Comput. Sci.* **1**, 542–549 (2021).
- Wendland, D. et al. Coherent dimension reduction with integrated photonic circuits exploiting tailored disorder. *J. Opt. Soc. Am. B* **40**, B35–B40 (2023).
- Wang, X. et al. Integrated photonic encoder for terapixel image processing. Preprint at *arXiv* <https://doi.org/10.48550/arXiv.2306.04554> (2023).
- Johnson, W. B. & Lindenstrauss, J. Extensions of Lipschitz mappings into a hilbert space. *Contemp. Math* **26**, 189–206 (1984).
- Gigan, S. Imaging and computing with disorder. *Nat. Phys.* **18**, 980–985 (2022).
- Mermelstein, M. D., Posey, R., Johnson, G. A. & Vohra, S. T. Rayleigh scattering optical frequency correlation in a single-mode optical fiber. *Opt. Lett.* **26**, 58–60 (2001).
- Sarma, R., Yamilov, A., Neupane, P., Shapiro, B. & Cao, H. Probing long-range intensity correlations inside disordered photonic nanostructures. *Phys. Rev. B - Condens. Matter Mater. Phys.* **90**, 014203 (2014).
- M. Scholz and R. Vigarío, “Nonlinear PCA: A new hierarchical approach,” *Proc. 10th Eur. Symp. Artif. Neural Networks*, 439–444 (2002).
- Mika, S. et al. Kernel PCA and de-noising in feature spaces. *Adv. Neural Inf. Process. Syst.* **11**, 536–542 (1999).
- Lever, J., Krzywinski, M. & Altman, N. Principal component analysis. *Nat. Methods* **14**, 641–642 (2017).
- Cortes, C. & Vapnik, V. Support-vector networks. *Mach. Learn.* **20**, 273–297 (1995).
- Noble, W. S. What is a support vector machine?. *Comput. Biol* **24**, 1565–1567 (2006).
- Gorman, R. P. & Sejnowski, T. J. Analysis of hidden units in a layered network trained to classify sonar targets. *Neural Networks* **1**, 75–89 (1988).
- Huang, G., Zhu, Q. & Siew, C. Extreme learning machine: theory and applications. *Neurocomputing* **70**, 489–501 (2006).
- Bin Huang, G., Zhou, H., Ding, X. & Zhang, R. Extreme learning machine for regression and multiclass classification. *IEEE Trans. Syst. Man Cybern. Part B Cybern.* **42**, 513–529 (2012).
- Pierangeli, D., Marcucci, G. & Conti, C. Photonic extreme learning machine by free-space optical propagation. *Photonics Res.* **9**, 1446–1454 (2021).
- LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2323 (1998).
- Miscuglio, M. et al. Massively parallel amplitude-only Fourier neural network. *Optica* **7**, 1812 (2020).
- Valley, G. C. Photonic analog-to-digital converters. *Opt. Express* **15**, 1955–1982 (2007).
- Corning Incorporated. <https://www.corning.com/media/worldwide/coc/documents/Fiber/PI-1463-AEN.pdf>.
- Miller, D. A. B. Energy consumption in optical modulators for interconnects. *Opt. Express* **20**, A293–A308 (2012).
- Wada, O. et al. 5 GHz-band CMOS direct digital RF modulator using current-mode DAC. *Asia-Pacific Microw. Conf. Proc. APMC*, 1118–1120 (IEEE, 2012).
- Caragiulo, B. M. P., Daigle, C. <https://github.com/pietro-caragiulo/survey-DAC>. *DAC Perform. Surv. 1996–2020*.
- Murmann, B. <https://github.com/bmurmann/ADC-survey>. *ADC Perform. Surv. 1997–2023*.
- Nozaki, K. et al. Femtofarad optoelectronic integration demonstrating energy-saving signal conversion and nonlinear functions. *Nat. Photonics* **13**, 454–459 (2019).

53. Li, G. et al. 25Gb/s 1V-driving CMOS ring modulator with integrated thermal tuning. *Opt. Express* **19**, 20435–20443 (2011).
54. V. A. Handerek, et al. Improved optical power budget in distributed acoustic sensing using enhanced scattering optical fibre. *26th Int. Conf. Opt. Fiber Sensors*, TuC5 (2018).
55. Guo, H., Liu, F., Yuan, Y., Yu, H. & Yang, M. Ultra-weak FBG and its refractive index distribution in the drawing optical fiber. *Opt. Express* **23**, 4829–4838 (2015).
56. Redding, B. et al. Low-noise distributed acoustic sensing using enhanced backscattering fiber with ultra-low-loss point reflectors. *Opt. Express* **28**, 14638–14647 (2020).
57. Masoudi, A. & Newson, T. P. Contributed review: distributed optical fibre dynamic strain sensing. *Rev. Sci. Instrum.* **87**, 011501 (2016).
58. Murray, M. J., Murray, J. B., Schermer, R. T., McKinney, J. D. & Redding, B. High-speed RF spectral analysis using a Rayleigh backscattering speckle spectrometer. *Opt. Express* **31**, 20651–20664 (2023).
59. Pour Fard, M. M. et al. Experimental realization of arbitrary activation functions for optical neural networks. *Opt. Express* **28**, 12138–12148 (2020).
60. Ashtiani, F., Geers, A. J. & Aflatouni, F. An on-chip photonic deep neural network for image classification. *Nature* **606**, 501–506 (2022).
61. Alem, M., Soto, M. A. & Thévenaz, L. Analytical model and experimental verification of the critical power for modulation instability in optical fibers. *Opt. Express* **23**, 29514–29532 (2015).
62. Sze, V., Chen, Y. H., Yang, T. J. & Emer, J. S. Efficient processing of deep neural networks: a tutorial and survey. *Proc. IEEE* **105**, 2295–2329 (2017).
63. Murray, M. J., Murray, J. B., Ogden, H. M. & Redding, B. Dynamic temperature-strain discrimination using a hybrid distributed fiber sensor based on Brillouin and Rayleigh scattering. *Opt. Express* **31**, 287–300 (2023).

Acknowledgements

B.R. and J.M. acknowledge support from the U.S. Naval Research Lab (6.1 Base Program). R.S. acknowledges the support by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-003525. This paper describes objective technical results and analysis. Any subjective views or

opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

Author contributions

B.R., J.B.M., J.D.H., S.S.P., Z.Z., and R.S. conceived of the experiments. B.R. and J.B.M. performed the experiment. B.R., J.B.M., J.D.H., S.S.P., and R.S. analyzed the data. B.R., Z.Z., and S.S.P. performed the energy consumption analysis. B.R. wrote the paper with contributions from all authors.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Brandon Redding.

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

Reprints and permissions information is available at

<http://www.nature.com/reprints>

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

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons 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/>.

This is a U.S. Government work and not under copyright protection in the US; foreign copyright protection may apply 2024