



Large-scale photonic natural language processing

CARLO M. VALENSISE,¹  IVANA GRECCO,² DAVIDE PIERANGELI,^{1,2,3,*} AND CLAUDIO CONTI^{1,2,3} 

¹Enrico Fermi Research Center (CREF), 00184 Rome, Italy

²Physics Department, Sapienza University of Rome, 00185 Rome, Italy

³Institute for Complex Systems, National Research Council (ISC-CNR), 00185 Rome, Italy

*Corresponding author: davide.pierangeli@roma1.infn.it

Received 10 August 2022; revised 27 September 2022; accepted 8 October 2022; posted 11 October 2022 (Doc. ID 472932); published 28 November 2022

Modern machine-learning applications require huge artificial networks demanding computational power and memory. Light-based platforms promise ultrafast and energy-efficient hardware, which may help realize next-generation data processing devices. However, current photonic networks are limited by the number of input-output nodes that can be processed in a single shot. This restricted network capacity prevents their application to relevant large-scale problems such as natural language processing. Here, we realize a photonic processor for supervised learning with a capacity exceeding 1.5×10^{10} optical nodes, more than one order of magnitude larger than any previous implementation, which enables photonic large-scale text encoding and classification. By exploiting the full three-dimensional structure of the optical field propagating in free space, we overcome the interpolation threshold and reach the over-parameterized region of machine learning, a condition that allows high-performance sentiment analysis with a minimal fraction of training points. Our results provide a novel solution to scale up light-driven computing and open the route to photonic natural language processing. © 2022

Chinese Laser Press

<https://doi.org/10.1364/PRJ.472932>

1. INTRODUCTION

Advanced artificial intelligence systems are becoming extremely demanding in terms of training time and energy consumption [1,2]. An ever-larger number of trainable parameters is required to exploit over-parameterization and achieve state-of-the-art performances [3,4]. Large-scale, energy-efficient computational hardware is becoming a subject of intense interest.

Photonic neuromorphic computing systems [5–7] offer large throughput [8] and energetically efficient [9] hardware accelerators based on integrated silicon photonic circuits [10–14], engineered meta-materials [15], or 3D-printed diffractive masks [16,17]. Other light-based computing architectures leverage reservoir computing (RC) [18] and extreme learning machine (ELM) [19] computational paradigms, where the input data is mapped into a feature space through a fixed set of random weights and training is performed only on the linear readout layer. Optical reservoir computers [20–33] and photonic extreme learning machines (PELMs) [34–37] apply successfully to various learning tasks, ranging from time series prediction [38,39] to image classification [40,41]. Despite the remarkable performances achieved by these architectures, their impact on large-scale problems is limited, mainly due to size constraints.

Here, we demonstrate photonic machine learning at an ultralarge scale by mapping the optical output layer in the full three-dimensional (3D) structure of the optical field. This original approach offers inherent scalability to the optical device and effortless access to the over-parameterized learning region. The 3D-PELM that we implement can process simultaneously up to 250,000 total input-output nodes via spatial light modulation, featuring a total network capacity one order of magnitude larger than existing optical reservoir computers for supervised learning. Our large-scale photonic network allows us to optically implement a massive text classification problem, the sentiment analysis of the Internet Movie Database (IMDb) [44], and enable, for the first time to our knowledge, the observation of the double descent phenomenon [45–47] on a photonic computing device. We demonstrate that, thanks to its huge number of optical output nodes, the 3D-PELM successfully classifies text by using only a limited number of training points, realizing energy-efficient, large-scale text processing.

2. RESULTS

A. Three-Dimensional PELM

A PELM classifies a dataset with N points $\mathbf{X} = [X_1, \dots, X_N]$ by mapping the input sample $X_i \in \mathbb{R}^L$ into a high-dimensional feature space through a nonlinear transformation that is

performed by a photonic platform [34]. To perform the classification, the output matrix \mathbf{H} , which contains the measured optical signals, is linearly combined with a set of M trainable readout weights β (see Appendix A). In all previous PELM realizations [25,34,35,39,41], the intensities stored in \mathbf{H} only contain information on the optical field at a single spatial or temporal plane. Here, to scale up the optical network, we exploit the entire three-dimensional optical field. Figure 1(B) shows a schematic of our 3D-PELM: the photonic network uses as output nodes the full 3D structure of the optical field propagating in free space. The transformation of the input data sample into an optical intensity volume occurs through linear wave mixing by coherent free-space propagation and intensity detection. Specifically, the 3D optical mapping reads as

$$\mathbf{H}_{3D} = \sum_j \mathbf{H}_j, \quad (1)$$

$$\mathbf{H}_j = G_j[\mathbf{M}_j \exp i(\mathbf{X} + \mathbf{W})], \quad (2)$$

where \mathbf{M}_j is the transfer matrix that models coherent light propagation to the j th detection plane, e.g., the discretized diffraction operator. The \mathbf{M}_j has complex coefficients, and thus the photonic scheme implements a complex-valued neural network [42,43]. In Eq. (2) the input dataset is encoded by phase modulation, while \mathbf{W} is a pre-selected random matrix that embeds the input signal and provides network biases. G_j is the response function of the j th detector to the impinging light field.

In experiments, to collect uncorrelated output data \mathbf{H}_j , we employ multiple cameras that simultaneously image the intensity in distinct far-field planes (see Appendix B). These speckle-like distributions [Fig. 1(C)] are uncorrelated to each other to acquire non-redundant information. Importantly, exploiting

the intensity of multiple optical planes is crucial for increasing the effective number of readout channels. In fact, the intensity on a single camera has a finite correlation length and near pixels are necessarily correlated, which makes the number of independent channels lower than the available sensor pixels. Our 3D scheme circumvents this limitation and allows to achieve a huge number of output nodes.

B. Optical Encoding of Natural Language Sentences

The natural language processing (NLP) task we considered is the classification of the IMDb dataset, constructed by Maas *et al.* [44]. To encode natural language on the optical setup, we need to develop an optics-specific representation of the text data. Since text classification tasks have remained elusive in photonic neuromorphic computing, we here introduce the problem. A similar issue is crucial also in digital NLP, where the text encoding problem consists in obtaining a numerical representation of the input that is convenient for digital learning architectures [48]. Specifically, given a vocabulary, i.e., the set of all unique words in the corpus, a basic encoding method is the one-hot technique, in which each word corresponds to a Boolean vector with vocabulary size V . Paragraphs of different lengths can be represented by Boolean vectors of size V , with the element $x_{ki} = 1$, which indicates the presence of the k th vocabulary word in the paragraph X_i [Fig. 1(A)]. An input dataset with N samples, in which each database entry is a text paragraph [see Fig. 1(A)], thus becomes an $N \times V$ matrix. The dimension of the single entry X_i scales with the number of words V in the vocabulary ($10^4 - 10^5$). Photonic text processing hence necessitates an optical platform capable of encoding large-size input data. Given the large number of input modes L supported by spatial light modulation, our 3D-PELM is the most convenient scheme for the scope.

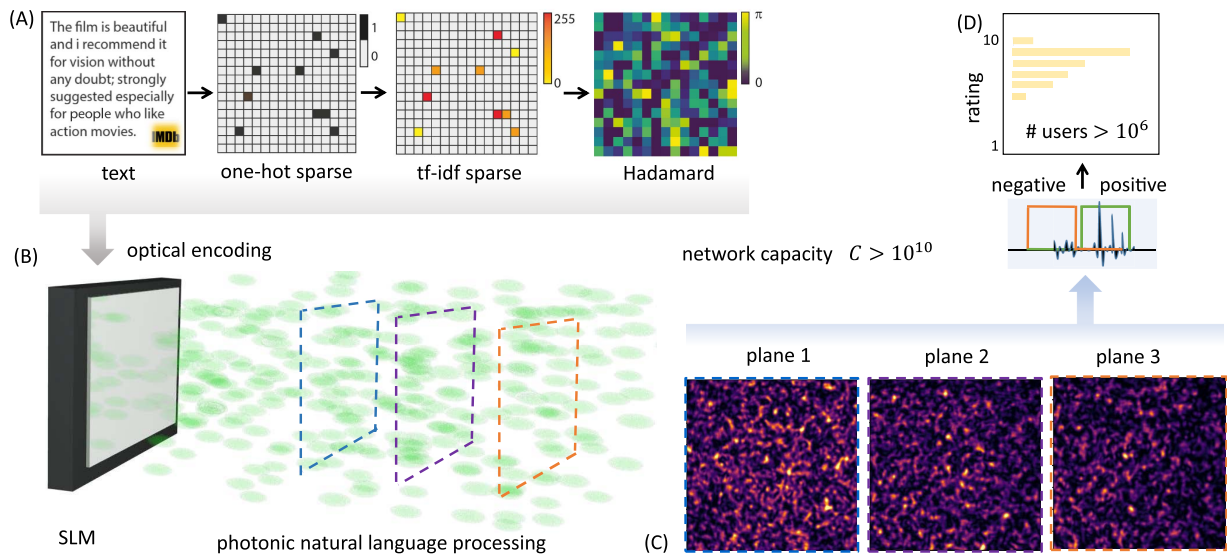


Fig. 1. Three-dimensional PELM for language processing. (A) The text database entry is a paragraph of variable length. Text pre-processing: a sparse representation of the input paragraph is mapped into a Hadamard matrix with phase values in $[0, \pi]$. (B) The mask is encoded into the optical wavefront by a phase-only SLM. Free-space propagation of the optical field maps the input data into a 3D intensity distribution (speckle-like volume). (C) Sampling the propagating laser beam in multiple far-field planes enables upscaling the feature space. Intensities picked from all the spatial modes form the output layer \mathbf{H}_{3D} that undergoes training via ridge regression. By using three planes ($j = 3$), we get a network capacity $C > 10^{10}$. (D) The example shows a binary text classification problem for large-scale rating.

Optical encoding of natural language in the 3D-PELM requires using a spatial light modulator (SLM) with a fixed number of input modes for sentences with variable lengths. We follow the one-hot method employed in digital NLP. We use an extension of one-hot encoding, the so-called tf-idf representation, and we construct the input representation $\mathbf{X}_{\text{tfidf}}$. In information retrieval, the tf-idf(i, j) statistics reflects how important a word k is to a document j (frequency n_{kj}) in a collection of documents. It is defined as $\text{tf-idf}(k, j) = \frac{n_{kj}}{V} \cdot \log_{10}(\frac{N}{S_k})$, where V is the length of a paragraph and S_k is the number of sentences containing the k th word. Within the tf-idf representation, each paragraph becomes a sparse, real vector [see Fig. 1(A)] whose non-zero elements are the tf-idf values of the words composing it. To optically encode these sparse data, we applied to $\mathbf{X}_{\text{tfidf}}$ the Walsh–Hadamard transform (WHT) [49], which decomposes a discrete function in a superposition of Walsh functions. The transformed dataset \mathbf{X}_{WHT} is a dense matrix [Fig. 1(A)], which is displayed on the SLM as a phase mask within the $[0, \pi]$ range [$\mathbf{X} = \mathbf{X}_{\text{WHT}}$ in Eq. (2)]. By exploiting the large number of available input pixels, we encode paragraphs with a massive size V containing hundreds of thousands of words.

C. Observation of the Photonic Double Descent

Machine learning models with a very large number of trainable parameters show unique features with respect to smaller models, such as the double descent phenomenon. This effect is a resonance in the neural network performance, ruled by the ratio between the number of trainable parameters M and the

number of training points N_{train} . In the under-parameterized regime ($M < N_{\text{train}}$), good performances are obtained by balancing model bias and variance. As M grows, models tend to overfit training data, and prediction performance gets worse until the so-called interpolation threshold ($M = N_{\text{train}}$) is reached. At this point the model optimally interpolates the training data and the prediction error is maximum. Beyond this resonance, in the over-parameterized regime, the model keeps interpolating training points, but performances on the test set reach the global optimum [50].

We experimentally implement photonic NLP and investigate the double descent effect on our large-scale 3D-PELM. Specifically, we analyze the classification accuracy for the IMDb task (Appendix D) as a function of the features M and training set size N_{train} . In Fig. 2(A) we report the observation of the double descent. We observe a dip in test accuracy as the number of channels M reaches the number of training points N_{train} . Beyond this resonance (interpolation threshold), we find the over-parameterized region in which maximum accuracy on the training set is achieved. The behavior is obtained via training on $N_{\text{train}} \simeq 1.2 \times 10^4$ examples and using fixed train/test split ratio of 0.67/0.33. In this case, the larger classification accuracy is found in the under-parameterized region (0.77). Conversely, in Fig. 2(B), we consider a much smaller number of training points, $N_{\text{train}} \simeq 10^3$. Remarkably, we reach the same optimal accuracy despite using only a fraction of the available training points. This observation reveals a particularly favorable learning region of the 3D-PELM in which only a reduced number of training points are required to reach high

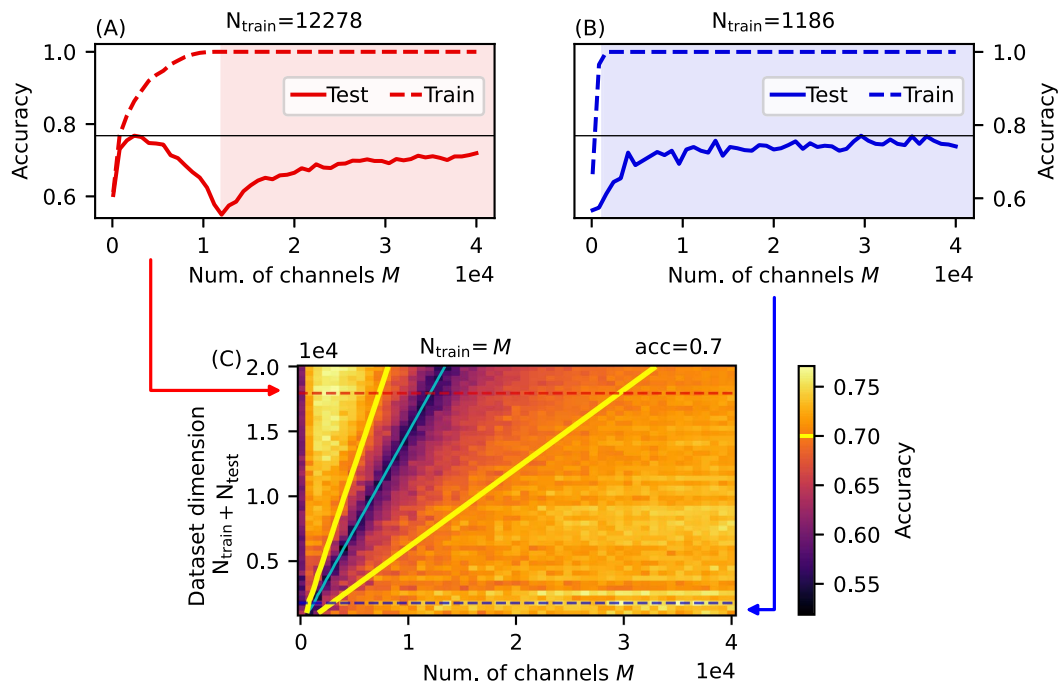


Fig. 2. Photonic sentiment analysis. (A), (B) Training and test accuracy of the 3D-PELM on the IMDb dataset as a function of the number of output channels. The shaded area corresponds to the over-parameterized region. The configuration in (B) allows us to reach very high accuracy in the over-parameterized region with a dataset limited to $N_{\text{train}} = 1186$ training points. In (A), the same accuracy is reached in the under-parameterized region with $N_{\text{train}} = 12,278$. Black horizontal lines correspond to the maximum test accuracy achieved (0.77). (C) IMDb classification accuracy by varying the number of features M and training dataset size N_{train} . The boundary between the under and over-parameterized region (interpolation threshold), $N_{\text{train}} = M$, is characterized by a sharp accuracy drop (cyan contour line).

accuracy levels. From the operational standpoint, it is much more efficient to measure in parallel a large number of modes rather than sequentially processing many training examples. In Fig. 2(C) we report the full dynamics of the double descent phenomenon by continuously varying both M and N_{train} . We observe that the accuracy dip shifts with a constant velocity. This result shows the existence of an optimal learning region that is accessible on the 3D-PELM thanks to its large number of photonic nodes.

D. Sentiment Analysis at Ultralarge Scale

The observed operational advantage of the over-parameterized region indicates that, by further increasing the number of output modes, one can increase training effectiveness and/or performances [3]. In our 3D-PELM we reached $M = 120,000$ readout channels that are independent of each other. We also considered larger input spaces, extending the vocabulary to include the whole set of words in the corpus ($V \simeq 2^{17}$). Figure 3(A) shows that the test accuracy reaches a plateau as M increases in the over-parameterized region. Saturation indicates that all the essential information encoded within the optical field has been extracted through the available channels. However, we observe a change in the performance as we employ more input features L [Figs. 3(B) and 3(C)]. To estimate the rate by which performance improves as more channels were used for training, we estimate the angular coefficient m of the linear growth that precedes the plateau. Although the onset of saturation can be estimated using diverse criteria, and m varies depending on the measured range, its trend when increasing L remains unaltered. We note a relevant enhancement of m for $L = 362 \times 362$, which indicates that PELMs featuring larger input spaces are able to reach optimal performances with a lower number of parameters. Figure 3(D) reports the accuracy as a function of the training-test split ratio (keeping fixed $N_{\text{train}} = 6.7 \times 10^3$) for $M = 0.8 \times 10^5$ and $M = 1.2 \times 10^5$.

Importantly, a limited number of examples ($\sim 20\%$) are enough to reach maximum accuracy in the over-parameterized region.

E. Optical Network Capacity

To establish a comparison among the various photonic neuromorphic computing devices, we introduce the optical network capacity as a useful ingredient related to the over-parameterization context. We define the capacity \mathcal{C} of a generic optical neural network as the product between the number of input and output nodes that can be processed in a single iteration, $\mathcal{C} = L \times M$. This quantity gives direct information on the kind of large-scale problems that can be implemented on the optical setup. It depends only on the number of controllable nodes, which is the quantity that is believed to play the main role in big data processing [3]. Specifically, L sets the size of the dataset that can be encoded, while M reflects the embedding capacity of the network, as larger datasets necessitate larger feature spaces to be learned. Moreover, \mathcal{C} also furnishes an indication on how far the over-parameterized regime is for a given task. Useful over-parameterized conditions can be reached only if $\mathcal{C} \gg N \times L$. We remark that the capacity is not a measure of the processor accuracy. It is instead a useful quantity to compare the scalability of different photonic computing devices.

In Table 1 we report the optical network capacity for various photonic processors that have been recently demonstrated. We focus on photonic platforms that exploit RC and ELM paradigms, since these devices are suitable for big data processing and may present the so-called photonic advantage at a large scale [51]. Our 3D-PELM has a record capacity $\mathcal{C} = 362 \times 362 \times 120,000 \simeq 1.5 \times 10^{10}$, more than one order of magnitude larger than any other optical processor that has been demonstrated on supervised learning. Moreover, while increasing the capacity is challenging on many optical reservoir computers, our 3D-PELM can be scaled up further by enlarging the measured volume of the optical field. For example, the optical

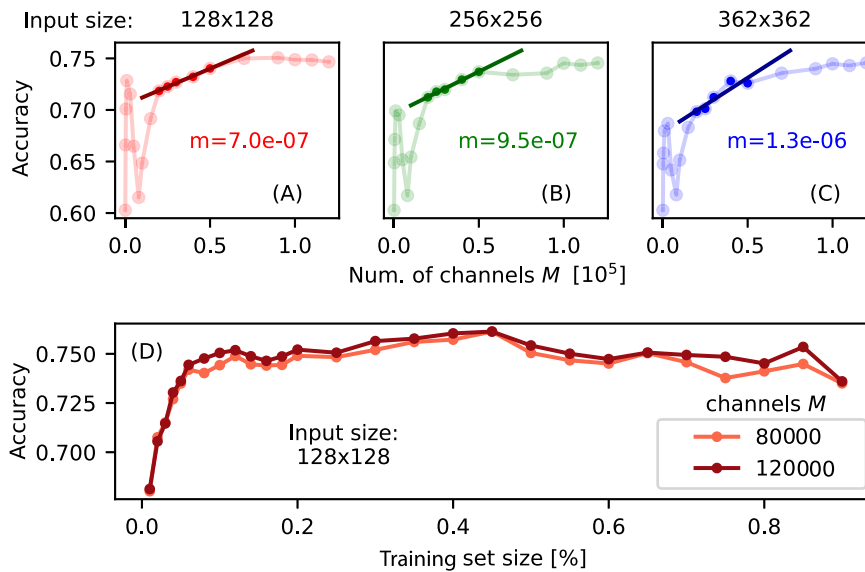


Fig. 3. Performances at ultralarge scale. (A)–(C) Test accuracy as a function of M for different input sizes L . In all cases, the 3D-PELM performance saturates in the over-parameterized region, reaching a plateau. A linear fit of the data preceding the plateau shows that the onset of the saturation is faster for datasets with a larger input space. The corresponding angular coefficient m is inset in each panel. (D) Test accuracy varying the training set size for $M = 0.8 \times 10^5$ and $M = 1.2 \times 10^5$.

Table 1. Maximum Network Capacity of Current Photonic Neuromorphic Computing Hardware for Supervised Learning

Working Principle	M	L	C	Machine Learning Task	Ref.
Time-multiplexed cavity	1400	7129	10^7	Regression	[39]
Amplitude modulation	16,384	2000	10^8	Human action recognition	[27]
Frequency multiplexing	200	640	10^5	Time series recovery	[41]
Optical multiple scattering	50,000	64	10^6	Chaotic series prediction	[38]
Amplitude Fourier filtering	1024	43,263	10^7	Image classification	[30]
Multimode fiber	240	240	10^5	Classification, regression	[35]
Free-space propagation	6400	784	10^6	Classification, regression	[34]
3D optical field	120,000	131,044	10^{10}	Natural language processing	3D-PELM

processor in Ref. [52], which is exploited for randomized numerical linear algebra, suggests that our scheme can reach a larger capacity by exploiting large-area cameras and multiple optical scattering.

F. Comparison with Digital Machine Learning

To validate further the capability of our photonic setup for text processing, we compare the accuracy of our 3D-PELM with various general-purpose digital neural networks on the IMDb task. To underline the overall impact of the over-parameterization on the performance, we consider two opposite parameterization regimes, $M = 1 \times 10^3$ and $M = 4 \times 10^4$, and two distinct conditions for the dataset size, $N_{\text{train}} = 6700$ [Fig. 4(A)] and $N_{\text{train}} = 1500$ [Fig. 4(B)]. Since the

learning principle of the 3D-PELM is based on kernel methods (see Ref. [34]), we implement a support vector machine (SVM) and a ridge regression based on nonlinear random projections (RP) [25] as representative models of the kernel algorithms class. The input for the SVM and the RP is the tf-idf-encoded dataset, but no significant differences are found using the Hadamard dataset. We also simulate the 3D-PELM, i.e., we evaluate our photonic scheme by using Eq. (2) to generate the output data. A convolutional neural network (CNN) with the same number of trainable weights is used as an additional benchmark model. Details on the various digital models are reported in Appendix E. The results for a split ratio $N_{\text{train}}/N_{\text{test}} = 0.67/0.33$ are shown in Fig. 4. Overall, we observe that the photonic device performs sentiment analysis on the IMDb with an accuracy comparable with standard digital methods. The SVM sets the maximum expected accuracy (0.86). The device over-parameterized accuracy (0.75 for $N_{\text{train}} = 6700$) is mainly limited by the limited precision of the optical components and noise, and it could be enhanced up to 0.83 (3D-PELM numerics) by using a 64-bit camera. In fact, when operating with 8-bit precision, both the simulated 3D-PELM and the RP models achieve performance that agrees well with the experiments. Interestingly, Fig. 4(B) indicates that, for a limited set of training samples, the 3D-PELM surpasses the SVM and also the CNN. This points out an additional advantage that the photonic setup achieves thanks to over-parameterization. Not only accuracy is improved with respect to the standard under-parameterized regime, but the device operates effectively in conditions where digital models are less accurate.

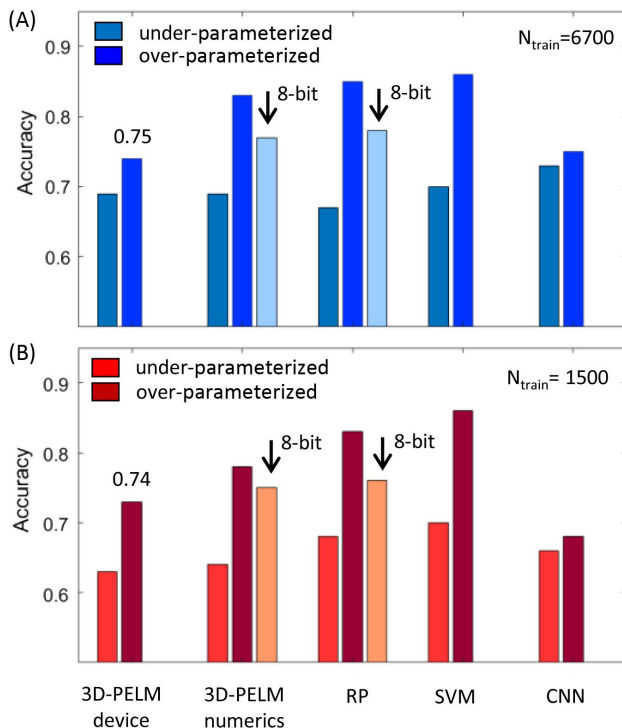


Fig. 4. Analysis of the IMDb accuracy. (A), (B) The comparison reports the accuracy for the experimental device (3D-PELM device), the simulated device (3D-PELM numerics), the random projection method with ridge regression (RP), the support vector machine (SVM), and a convolutional neural network (CNN) in both the under-parameterized ($M = 1 \times 10^3$) and over-parameterized ($M = 4 \times 10^4$) regimes, for (A) $N_{\text{train}} = 6700$ and (B) $N_{\text{train}} = 1500$. 8-bit numerical results, when applicable, refer to the over-parameterized regime.

3. DISCUSSION

We have reported the first, to our knowledge, photonic implementation of natural language processing by performing sentiment analysis on the IMDb dataset. Our results demonstrate the feasibility of modern large-scale computing tasks in photonic hardware. They can be potentially improved both in terms of performance, via low-noise high-resolution cameras and more complex encoding/training algorithms, and applications. Further developments include the implementation of advanced tasks such as multiclass sentiment analysis, which allows discriminating positive and negative sentences into different levels. Moreover, using different datasets such as the Stanford Sentiment Treebank (SST-5), we could train our 3D-PELM for providing ratings. On the other hand, our versatile optical setting may allow implementation of alternative strategies for photonic NLP, including dimensionality

reduction of the input data via optical random projection [53]. An optical scheme that preserves the sequential nature of language may be also conceived by introducing recurrent mechanisms in our 3D-PELM. Another interesting direction is developing specific encoding schemes for representing text in the optical domain. Viable approaches include adapting bidimensional text representations such as quick response (QR) codes. Furthermore, by following the concept of word embedding used in digital NLP [54], photonic hardware may be used to learn a purely optical word embedding, including semantics, word ordering, and syntax.

In conclusion, modern machine learning architectures rely on over-parameterization to reach the global optimum of objective functions. In photonic computing, achieving over-parameterization for large-scale problems is challenging. Here, we realize a novel photonic computing device that, thanks to its huge network capacity, naturally operates in the over-parameterized region. The 3D-PELM is a very promising solution for developing a fully scalable photonic machine-learning platform. Our photonic processor enables the first link between NLP and optical computing hardware, opening a new vibrant research stream toward fast and energy-efficient artificial intelligence tools.

APPENDIX A: PELM FRAMEWORK

In ELMs, a dataset with N data points $\mathbf{X} = [(X_1, y_1), \dots, (X_N, y_N)]$, with $(X_i, y_i) \in \mathbb{R}^L \times \mathbb{R}$, is mapped into a higher-dimensional feature space through a nonlinear transformation $g(\mathbf{X})$, yielding the hidden-layer output matrix $\mathbf{H} = [g(X_1), \dots, g(X_N)]$, with $g(X_i) \in \mathbb{R}^M$. A set of weights β is learned via ridge regression such that the linear combination $\tilde{\mathbf{Y}} = \mathbf{H}\beta$ well approximates the true label vector $\mathbf{Y} = [y_1, \dots, y_N]$. An explicit solution is $\beta = (\mathbf{H}^T \mathbf{H} + \lambda \mathbf{I})^{-1} \mathbf{H}^T \mathbf{Y}$, where λ is the regularization parameter and \mathbf{I} is the identity. The use of ridge regression allows reducing at the minimum training costs, which is crucial in applications requiring fast and reconfigurable learning. In a free-space PELM, the nonlinear mapping of input data is realized by a combination of free-space optical propagation and intensity detection [34]. The speckle-like pattern detected by a camera forms the nonlinear features of the input sample. The optical mapping reads as $\mathbf{H} = G[\mathbf{M} \exp i(\mathbf{X} + \mathbf{W})]$, where the input dataset \mathbf{X} is phase-encoded via spatial light modulation. \mathbf{W} is the embedding matrix, G is the detector response function, and \mathbf{M} is the discrete Fourier transform that models propagation to the focal plane. For more details on the PELM architecture and its working principle, see Ref. [34].

APPENDIX B: EXPERIMENTAL SETUP

The optical setup of the 3D-PELM is sketched in Fig. 1(B). A CW laser beam (532 nm, 100 mW) is expanded and collimated onto a phase-only SLM (Hamamatsu X13138, 1280 × 1024 pixels, 60 Hz, 8 bits), which encodes data on the optical wavefront. The SLM operates in the $[0, 2\pi]$ interval with a linear response function. Both the input and the embedding signal are encoded within the $[0, \pi]$ range and superimposed. After free-space propagation through a lens ($f = 150$ mm), the far-field three-dimensional intensity distribution is accessed

simultaneously by three separated imaging systems and detected by three cameras (Basler Ace 2 Pro, 1920 × 1200 pixels, 12 bits). The speckle-like field spatially decorrelates during propagation; the transverse intensity is transformed in a random way from one plane to the other, provided that these planes are at distances larger than the longitudinal field correlation length. To carry out the measurements, the propagating beam is divided by a series of beam splitters into separated imaging paths with different optical length. The far-field planes are selected in order to sample three speckle-like distributions having a negligible mutual intensity correlation. This is confirmed by intensity correlation measurements. Each additional plane contains additional information. In fact, keeping fixed M , we find an improvement in the classification accuracy when using three far-field planes instead of a single intensity distribution. In this condition, each camera maps a distinct portion of the high-dimensional feature space where the input field is projected. Importantly, speckle grains typically extend over adjacent pixels [Fig. 1(C)]. To avoid spatial correlations within the single speckle pattern, pixels are grouped into square modes (channels), and the average intensity on each channels forms the output signals. The intensity correlation length on the camera cannot be reduced to the single pixel due to the minimal size of the input modes that is limited by the SLM pixel pitch. For this reason, output channels cannot be increased by simply using all the sensor pixels, but multiple optical planes are necessary. The measured set of speckle-like patterns is elaborated on a conventional computer.

APPENDIX C: 3D-PELM TRAINING

Training operates by loading the randomly ordered input dataset on the SLM and measuring three speckle-like intensity distributions for each input sample (paragraph). From the acquired signals, we randomly select M of the available channels, and we use the corresponding M intensity values to form the output dataset. This dataset is split randomly into a training and a test datasets by using a split ratio $N_{\text{train}}/N_{\text{test}} = 0.67/0.33$, which is kept fixed throughout the analysis. All the classification accuracies we obtained refer to this hyperparameter. The training dataset is the output matrix \mathbf{H} used for training (see Appendix A). Since our task corresponds to a binary classification problem, the target labels are $y_i \in \{0, 1\}$. According to the PELM framework, training by ridge regression consists in solving the equation for the β vector. This matrix inversion is performed using the ML Python Library Sci-kit learn. The regularization parameter λ is varied in the $[10^{-4}, 10^2]$ range. As its value does not influence considerably the testing accuracy, we use $\lambda = 10^{-4}$. Given a predicted output \hat{y}_i , the predicted label is given by $\tilde{y}_i = \Theta(\hat{y}_i - 0.5)$, where Θ is the Heaviside step function. Text classification performances are evaluated via the accuracy $\sum_i [y_i = \tilde{y}_i]/N$, where $[\cdot]$ is the Iverson bracket defined as $[P] = 1$ if P is true and 0 otherwise.

APPENDIX D: NLP TASK, TEXT PRE-PROCESSING, AND OPTICAL ENCODING

We consider the IMDb dataset [44], a widely used benchmark for polarity-based text classification [55]. The dataset consists of 5×10^4 movie reviews, each one forming the i th input data

sample. Each review is made of a paragraph consisting of multiple sentences and expresses an overall positive or negative sentiment (target variable). This sentiment analysis task hence consists in a large-scale binary classification. To implement the problem, the whole set of IMDb sentences is pre-processed by removing special characters (e.g., html), lowercasing, removing stop words, and lemmatizing. The length of the resulting vector can be varied by neglecting words that appear in the corpus with a frequency lower than a given threshold. To compute the WHT and obtain a dense text representation, zero-padding is used to reach a vector length $d = \lceil \log_2 V \rceil$. For example, results in Fig. 2 refer to data with $V = 10^4$ that are padded to $L = 2^{14} = 16,384$ and reshaped as a 128×128 square matrix, while the full vocabulary ($V = 131,044$) requires zero-padding, and paragraphs are arranged on 362×362 input modes (Fig. 3). After WHT computation, the values are normalized in the $[0, \pi]$ range and loaded on the SLM. Each element of the input \mathbf{X}_{WHT} is displayed onto a square block of SLM pixels. The embedding matrix \mathbf{W} is a uniformly distributed random matrix for all the experiments. Its values w_{ij} are selected in $[0, \pi]$ before the training and kept fixed for a given photonic computation.

APPENDIX E: DIGITAL MACHINE LEARNING

The 3D-PELM is numerically simulated by following the optical mapping model in Eq. (1) (see also Ref. [34]). Data encoding and training operates in the same way as the photonic device, and the output matrix is computed by using Eq. (2) with $G(z) = |z|^2$ (quadratic nonlinearity). In the RP model, data projection in the feature space is obtained by matrix multiplication with a real-valued uniformly distributed random matrix and using a quadratic activation function. Training occurs by ridge regression. For simulating the 8-bit machines, the input data and the output values forming the linear readout layer are properly discretized to values with 8-bit precision. The SVM is a binary classifier, and thus it is especially suited for the IMDb task. We use the Sci-kit learn built-in function on the tf-idf-processed dataset, and the over-parameterized regime is evaluated by using $M = 10^3$. The CNN is composed of four convolutional layers with nonlinear activation and a fully connected layer. The CNN scheme is designed to have M total nodes, and its accuracy does not depend significantly on the network details and hyperparameters.

Funding. Museo Storico della Fisica e Centro Studi e Ricerche Enrico Fermi; Ministero dell'Università e della Ricerca (PRIN No. 20177PSCKT).

Acknowledgment. We thank I. MD Deen and F. Farrelly for technical support in the laboratory. D.P. and C. C. conceived the research. C.M.V. proposed application to NLP. D.P. developed the 3D photonic network. I.G. and D.P. carried out experimental measurements. C.M.V. and I. G. performed data analysis. D.P. and C.C. co-supervised the project. All authors contributed to the manuscript.

Disclosures. The authors declare no conflicts of interest.

Data Availability. The data that support the results presented in this study are available from the corresponding author upon reasonable request.

Code Availability. The codes developed within this research are available from the corresponding author upon reasonable request.

REFERENCES

1. D. Narayanan, M. Shoeybi, J. Casper, P. LeGresley, M. Patwary, V. A. Korthikanti, D. Vainbrand, P. Kashinkunti, J. Bernauer, B. Catanzaro, A. Phanishayee, and M. Zaharia, "Efficient large-scale language model training on GPU clusters using megatron-LM," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (ACM, 2021)*, pp. 1–15.
2. E. Strubell, A. Ganesh, and A. McCallum, "Energy and policy considerations for modern deep learning research," *Proc. AAAI Conf. Artif. Intell.* **34**, 13693–13696 (2020).
3. A. Chatelain, A. Djeghri, D. Hesslow, J. Launay, and I. Poli, "Is the number of trainable parameters all that actually matters?" arXiv:2109.11928 (2021).
4. N. C. Thompson, K. Greenewald, K. Lee, and G. F. Manso, "The computational limits of deep learning," arXiv:2007.05558 (2020).
5. G. Wetzstein, A. Ozcan, S. Gigan, S. Fan, D. Englund, M. Soljačić, C. Denz, D. A. B. Miller, and D. Psaltis, "Inference in artificial intelligence with deep optics and photonics," *Nature* **588**, 39–47 (2020).
6. B. J. Shastri, A. N. Tait, T. F. de Lima, W. H. P. Pernice, H. Bhaskaran, C. D. Wright, and P. R. Prucnal, "Photonics for artificial intelligence and neuromorphic computing," *Nat. Photonics* **15**, 102–114 (2021).
7. H. Zhou, J. Dong, J. Cheng, W. Dong, C. Huang, Y. Shen, Q. Zhang, M. Gu, C. Qian, H. Chen, Z. Ruan, and X. Zhang, "Photonic matrix multiplication lights up photonic accelerator and beyond," *Light Sci. Appl.* **11**, 30 (2022).
8. X. Xu, M. Tan, B. Corcoran, J. Wu, A. Boes, T. G. Nguyen, S. T. Chu, B. E. Little, D. G. Hicks, R. Morandotti, A. Mitchell, and D. J. Moss, "11 TOPS photonic convolutional accelerator for optical neural networks," *Nature* **589**, 44–51 (2021).
9. T. Wang, S.-Y. Ma, L. G. Wright, T. Onodera, B. Richard, and P. L. McMahon, "An optical neural network using less than 1 photon per multiplication," *Nat. Commun.* **13**, 123 (2022).
10. Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund, and M. Soljačić, "Deep learning with coherent nanophotonic circuits," *Nat. Photonics* **11**, 441–446 (2017).
11. A. N. Tait, T. F. de Lima, E. Zhou, A. X. Wu, M.-A. Nahmias, B. J. Shastri, and P. R. Prucnal, "Neuromorphic photonic networks using silicon photonic weight banks," *Sci. Rep.* **7**, 7430 (2017).
12. J. Feldmann, N. Youngblood, C. D. Wright, H. Bhaskaran, and W. H. P. Pernice, "All-optical spiking neurosynaptic networks with self-learning capabilities," *Nature* **569**, 208–214 (2019).
13. F. Stelzer, A. Röhm, R. Vicente, I. Fischer, and S. Yanchuk, "Deep neural networks using a single neuron: folded-in-time architecture using feedback-modulated delay loops," *Nat. Commun.* **12**, 5164 (2021).
14. J. Feldmann, N. Youngblood, M. Karpov, H. Gehring, X. Li, M. Stappers, M. Le Gallo, X. Fu, A. Lukashchuk, A. S. Raja, J. Liu, C. D. Wright, A. Sebastian, T. J. Kippenberg, W. H. P. Pernice, and H. Bhaskaran, "Parallel convolutional processing using an integrated photonic tensor core," *Nature* **589**, 52–58 (2021).
15. N. Mohammadi Estakhri, B. Edwards, and N. Engheta, "Inverse-designed metastructures that solve equations," *Science* **363**, 1333–1338 (2019).
16. X. Lin, Y. Rivenson, N. T. Yardimci, M. Veli, Y. Luo, M. Jarrahi, and A. Ozcan, "All-optical machine learning using diffractive deep neural networks," *Science* **361**, 1004–1008 (2018).
17. T. Zhou, X. Lin, J. Wu, Y. Chen, H. Xie, Y. Li, J. Fan, H. Wu, L. Fang, and Q. Dai, "Large-scale neuromorphic optoelectronic computing with

- a reconfigurable diffractive processing unit," *Nat. Photonics* **15**, 367–373 (2021).
18. D. Verstraeten, B. Schrauwen, M. D'Haene, and D. Stroobandt, "An experimental unification of reservoir computing methods," *Neural Netw.* **20**, 391–403 (2007).
 19. G. B. Huang, H. Zhou, X. Ding, and R. Zhang, "Extreme learning machine for regression and multiclass classification," *IEEE Trans. Syst. Man Cybern. B* **42**, 513–529 (2012).
 20. G. Van der Sande, D. Brunner, and M. C. Soriano, "Advances in photonic reservoir computing," *Nanophotonics* **6**, 561–576 (2017).
 21. D. Brunner, M. C. Soriano, C. Mirasso, and I. Fischer, "Parallel photonic information processing at gigabyte per second data rates using transient states," *Nat. Commun.* **4**, 1364 (2013).
 22. Q. Vinckier, F. Duport, A. Smerieri, K. Vandoorne, P. Bienstman, M. Haelterman, and S. Massar, "High-performance photonic reservoir computer based on a coherently driven passive cavity," *Optica* **2**, 438–446 (2015).
 23. L. Larger, A. Baylón-Fuentes, R. Martinenghi, V. S. Udaltsov, Y. K. Chembu, and M. Jacquot, "High-speed photonic reservoir computing using a time-delay based architecture: million words per second classification," *Phys. Rev. X* **7**, 011015 (2017).
 24. D. Ballarini, A. Gianfrate, R. Panico, A. Opala, S. Ghosh, L. Dominici, V. Ardizzone, M. D. Giorgi, G. Lerario, G. Gigli, T. C. H. Liew, M. Matuszewski, and D. Sanvitto, "Polaritonic neuromorphic computing outperforms linear classifiers," *Nano Lett.* **20**, 3506–3512 (2020).
 25. A. Saade, F. Caltagirone, I. Carron, L. Daudet, A. Dremeau, S. Gigan, and F. Krzakala, "Random projections through multiple optical scattering: approximating kernels at the speed of light," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2016), pp. 6215–6219.
 26. J. Bueno, S. Maktoobi, L. Froehly, I. Fischer, M. Jacquot, L. Larger, and D. Brunner, "Reinforcement learning in a large-scale photonic recurrent neural network," *Optica* **5**, 756–760 (2018).
 27. P. Antonik, N. Marsal, D. Brunner, and D. Rontani, "Human action recognition with a large-scale brain-inspired photonic computer," *Nat. Mach. Intell.* **1**, 530–537 (2019).
 28. A. Röhm, L. Jaurigue, and K. Lüdge, "Reservoir computing using laser networks," *IEEE J. Sel. Top. Quantum Electron.* **26**, 7700108 (2020).
 29. U. Paudel, M. Luengo-Kovac, J. Pilawa, T. J. Shaw, and G. C. Valley, "Classification of time-domain waveforms using a speckle-based optical reservoir computer," *Opt. Express* **28**, 1225–1237 (2020).
 30. M. Miscuglio, Z. Hu, S. Li, J. George, R. Capanna, P. M. Bardet, P. Gupta, and V. J. Sorger, "Massively parallel amplitude-only Fourier neural network," *Optica* **7**, 1812–1819 (2020).
 31. S. Sunada and A. Uchida, "Photonic neural field on a silicon chip: large-scale, high-speed neuro-inspired computing and sensing," *Optica* **8**, 1388–1396 (2021).
 32. M. Borghi, S. Biasi, and L. Pavesi, "Reservoir computing based on a silicon microring and time multiplexing for binary and analog operations," *Sci. Rep.* **11**, 15642 (2021).
 33. X. Porte, A. Skalli, N. Haghighi, S. Reitzenstein, J. A. Lott, and D. Brunner, "A complete, parallel and autonomous photonic neural network in a semiconductor multimode laser," *J. Phys. Photon.* **3**, 024017 (2021).
 34. D. Pierangeli, G. Marcucci, and C. Conti, "Photonic extreme learning machine by free-space optical propagation," *Photon. Res.* **9**, 1446–1454 (2021).
 35. U. Teğin, M. Yıldırım, İ. Oğuz, C. Moser, and D. Psaltis, "Scalable optical learning operator," *Nat. Comput. Sci.* **1**, 542–549 (2021).
 36. D. Pierangeli, G. Marcucci, and C. Conti, "Neuromorphic computing device using optical shock waves," in *OSA Nonlinear Optics* (Optica Publishing Group, 2021), paper NTh1A-3.
 37. Z. Denis, I. Favero, and C. Ciuti, "Photonic Kernel machine learning for ultrafast spectral analysis," *Phys. Rev. Appl.* **17**, 034077 (2022).
 38. M. Rafayelyan, J. Dong, Y. Tan, F. Krzakala, and S. Gigan, "Large-scale optical reservoir computing for spatiotemporal chaotic systems prediction," *Phys. Rev. X* **10**, 041037 (2020).
 39. S. Ortín, M. C. Soriano, L. Pesquera, D. Brunner, D. San-Martín, I. Fischer, C. R. Mirasso, and J. M. Gutiérrez, "A unified framework for reservoir computing and extreme learning machines based on a single time-delayed neuron," *Sci. Rep.* **5**, 14945 (2015).
 40. R. Mirek, A. Opala, P. Comaron, M. Furman, M. Król, K. Tyszka, B. Seredyński, D. Ballarini, D. Sanvitto, T. C. H. Liew, W. Pacuski, J. Sufczyński, J. Szczytko, M. Matuszewski, and B. Piętko, "Neuromorphic binarized polariton networks," *Nano Lett.* **21**, 3715–3720 (2021).
 41. A. Lupo and S. Massar, "Parallel extreme learning machines based on frequency multiplexing," *Appl. Sci.* **12**, 214 (2021).
 42. G. Marcucci, D. Pierangeli, and C. Conti, "Theory of neuromorphic computing by waves: machine learning by rogue waves, dispersive shocks, and solitons," *Phys. Rev. Lett.* **125**, 093901 (2020).
 43. H. Zhang, M. Gu, X. D. Jiang, J. Thompson, H. Cai, S. Paesani, R. Santagati, A. Laing, Y. Zhang, M. H. Yung, Y. Z. Shi, F. K. Muhammad, G. Q. Lo, X. S. Luo, B. Dong, D. L. Kwong, L. C. Kwek, and A. Q. Liu, "An optical neural chip for implementing complex-valued neural network," *Nat. Commun.* **12**, 457 (2021).
 44. A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts, "Learning word vectors for sentiment analysis," in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies* (Association for Computational Linguistics, 2011), pp. 142–150.
 45. M. Belkin, D. Hsu, S. Ma, and S. Mandal, "Reconciling modern machine-learning practice and the classical bias–variance trade-off," *Proc. Natl. Acad. Sci. USA* **116**, 15849–15854 (2019).
 46. M. S. Advani, A. M. Saxe, and H. Sompolinsky, "High-dimensional dynamics of generalization error in neural networks," *Neural Netw.* **132**, 428–446 (2020).
 47. S. Mei and A. Montanari, "The generalization error of random features regression: precise asymptotics and double descent curve," *Commun. Pure Appl. Math.* **75**, 667–766 (2020).
 48. K. Babić, S. Martinčić-Ipšić, and A. Meštrović, "Survey of neural text representation models," *Information* **11**, 511 (2020).
 49. A. Ashrafi, "Walsh–Hadamard transforms: a review," in *Advances in Imaging and Electron Physics* (Elsevier, 2017), pp. 1–55.
 50. M. Soltanolkotabi, A. Javanmard, and J. D. Lee, "Theoretical insights into the optimization landscape of overparameterized shallow neural networks," *IEEE Trans. Inf. Theory* **65**, 742–769 (2019).
 51. D. Pierangeli, M. Rafayelyan, C. Conti, and S. Gigan, "Scalable spin-glass optical simulator," *Phys. Rev. Appl.* **15**, 034087 (2021).
 52. D. Hesslow, A. Cappelli, I. Carron, L. Daudet, R. Lafargue, K. Müller, R. Ohana, G. Pariente, and I. Poli, "Photonic co-processors in HPC: using LightOn OPU for randomized numerical linear algebra," arXiv:2104.14429 (2021).
 53. M.-A. Miri, "Integrated random projection and dimensionality reduction by propagating light in photonic lattices," *Opt. Lett.* **46**, 4936–4939 (2021).
 54. T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," arXiv:1301.3781 (2013).
 55. A. U. Rehman, A. K. Malik, B. Raza, and W. Ali, "A hybrid CNN-LSTM model for improving accuracy of movie reviews sentiment analysis," *Multimedia Tools Appl.* **78**, 26597 (2019).