



OPEN

Efficient shallow learning as an alternative to deep learning

Yuval Meir¹, Ofek Tevet¹, Yarden Tzach¹, Shiri Hodassman¹, Ronit D. Gross¹ & Ido Kanter^{1,2}✉

The realization of complex classification tasks requires training of deep learning (DL) architectures consisting of tens or even hundreds of convolutional and fully connected hidden layers, which is far from the reality of the human brain. According to the DL rationale, the first convolutional layer reveals localized patterns in the input and large-scale patterns in the following layers, until it reliably characterizes a class of inputs. Here, we demonstrate that with a fixed ratio between the depths of the first and second convolutional layers, the error rates of the generalized shallow LeNet architecture, consisting of only five layers, decay as a power law with the number of filters in the first convolutional layer. The extrapolation of this power law indicates that the generalized LeNet can achieve small error rates that were previously obtained for the CIFAR-10 database using DL architectures. A power law with a similar exponent also characterizes the generalized VGG-16 architecture. However, this results in a significantly increased number of operations required to achieve a given error rate with respect to LeNet. This power law phenomenon governs various generalized LeNet and VGG-16 architectures, hinting at its universal behavior and suggesting a quantitative hierarchical time–space complexity among machine learning architectures. Additionally, the conservation law along the convolutional layers, which is the square-root of their size times their depth, is found to asymptotically minimize error rates. The efficient shallow learning that is demonstrated in this study calls for further quantitative examination using various databases and architectures and its accelerated implementation using future dedicated hardware developments.

Traditionally, artificial neural networks have been derived from brain dynamics, where synaptic plasticity modifies the connection strength between two neurons in response to their relative activities^{1,2}. The earliest artificial neural network was the Perceptron^{3,4}, which was introduced approximately 65 years ago, consisting of a feedforward classifier with many inputs and a single Boolean output unit. The development of more structured feedforward architectures with numerous convolutional and fully connected hidden layers, which can be increased to hundreds^{5,6}, as well as the development of their non-local training techniques, such as backpropagation (BP)^{7,8}, are required to address solutions to complex and practical classification tasks. These are essential components of the current implementation of deep learning (DL) algorithms. The underlying rationality of DL algorithms is that the first convolutional layer is sensitive to the appearance of a given pattern or symmetry in limited areas of the input, whereas the subsequent convolutional layers are expected to reveal large-scale features characterizing a class of inputs^{9,10}.

In a supervised learning scenario, a feedforward step is initially performed, in which the distance between the current and desired outputs for a given input is computed using a given error function. The BP procedure is utilized in the next step, where weights are updated to locally minimize the error function^{7,11}. Graphic processing units (GPUs) are used to accelerate this time-consuming computational process of multiplying large matrices and vectors, and its use is repeated several times over the training set until a possible desired test error is achieved. Architectures with an increasing number of hidden layers enable learning to be efficiently optimized for complex classification tasks, which goes together with the advancement of powerful GPU technology.

However, the brain's architecture differs significantly from that of DL and consists of very few feedforward layers^{12–14}, only one of which approximates the convolutional wiring, mainly from the retinal input to the first hidden layer^{12,15}. The key question driving our research is whether learning non-trivial classification tasks using brain-inspired shallow feedforward networks can achieve the same error rates as DL, while potentially requiring less computational complexity. A positive answer will question the need for DL architectures and might direct the development of unique hardware for the efficient and fast implementation of shallow learning. Additionally, it will demonstrate how brain-inspired shallow learning has advanced computational capability with reduced complexity and energy consumption^{16,17}.

¹Department of Physics, Bar-Ilan University, 52900 Ramat-Gan, Israel. ²Gonda Interdisciplinary Brain Research Center, Bar-Ilan University, 52900 Ramat-Gan, Israel. ✉email: ido.kanter@biu.ac.il

Results

LeNet^{18,19}, a five-layer prototype of a shallow feedforward architecture, has two convolutional layers with max-pooling operations and three successive fully connected layers (Fig. 1A). The first and second convolutional layers have $d_1 = 6$ and $d_2 = 16$ filters, respectively, representing the depth of each layer, and their convolutional layer sizes after max-pooling, $m_i \times m_i$ (*height* \times *width*), are 14×14 and 5×5 , respectively. One can notice that

$$\frac{d_2}{d_1} = \frac{16}{6} \simeq \frac{14}{5} = \frac{m_1}{m_2} = 2.8, \quad (1)$$

which hints on the following conservation law along the convolutional layers

$$\text{depth}_i \times m_i = \text{constant}, \quad (2)$$

where m_i^2 and depth_i represent the i th convolutional layer size and the number of filters, respectively. We minimized the LeNet error rates for the CIFAR-10 database²⁰ as a function of d_1 while maintaining the ratio d_2/d_1 constant using the stochastic gradient descent (SGD) algorithm^{21,22} (Fig. 1B, Supplementary Information). The results indicate decaying of error rates, ϵ , with increasing d_1 as a power law

$$\epsilon(d_1) = \frac{A}{(d_1)^\rho}, \quad (3)$$

with an exponent $\rho \sim 0.41$, even for small d_1 . Although the error rate of the original LeNet, $d_1 = 6$, is $\epsilon \simeq 0.23$, it can be further minimized by increasing d_1 . Any small ϵ can be achieved on the test set using generalized LeNet, a shallow architecture, based on the power law extrapolation for large d_1 values. However, its minimization for a given large d_1 is a heavy computational task that requires an exhaustive search in a hyper-parameter space that its values vary among layers, with an increasing number of epochs and complex scheduling. For instance, preliminary results of an incomplete optimization for $d_1 = 27$ and $d_2 = 72$ using at least 500 epochs indicate $\epsilon \sim 0.137$, which is close to the expected result of the extrapolated power law (Fig. 1B).

The conservation law (Eq. 2) was found to govern the convolutional layer sizes of the original VGG-16 architecture, which consists of 16 layers²³, except for the fifth convolution set, where the number of filters is bounded by 512 (Fig. 2A with $d = 64$). The n th ($n \leq 4$) convolution set has $d \cdot 2^{n-1}$ filters, where the convolutional layer size is $\frac{m}{2^{n-1}}$ ($n \leq 5$). The minimization of ϵ for VGG-16 and the CIFAR-10 database (Fig. 2A with $m = 32$) as a function of d results in a power law with a similar exponent to LeNet (Fig. 1B), $\rho \sim 0.4$ (Fig. 2B, Supplementary Information). The results allude to the universal behavior of power-law scaling (Eq. 3) which is independent of the architecture details, where d is the number of filters in the first convolutional layer. Additionally, the exponent, ρ , does not necessarily increase with the number of convolutional or hidden layers. Interestingly, the standard VGG-16 network ($d = 64$ in Fig. 2A), with batch normalization but without dropouts, results in $\epsilon \sim 0.065$ (Supplementary Information), which is identical to the reported test error with significant dropouts²⁴. Hence, the advantage of dropouts in the minimization of ϵ might be questionable in this case.

A shallow network's ability to achieve any small ϵ , based on the extrapolation of the power-law scaling (Fig. 1B), is accompanied by a significant reduction in computational complexity per epoch compared with a DL architecture (Fig. 2A). Complexity is measured as the number of multiplication-add (MAdd) operations per input during a forward and BP step^{25,26}. It is calculated as a function of the number of filters, d_1 and d in Figs. 1 and 2, respectively (Fig. 3A, Supplementary Information). In both cases, the number of operations per step scale as a quadratic polynomial with the number of filters and are derived from the following argument: When the number of filters in a convolutional layer is doubled, its computational complexity increases by a factor of four because its consecutive convolutional layer is also doubled. The origin of the linear terms in the quadratic polynomials (Fig. 3A) is mainly attributed to the input size of the first fully connected layer, which increased

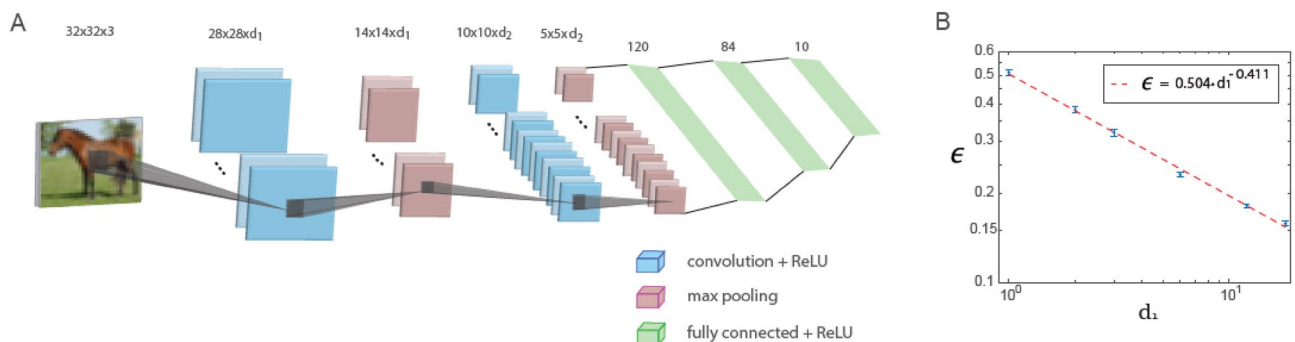


Figure 1. Learning in generalized LeNet architecture. (A) Generalized LeNet architecture for the CIFAR-10 database ($32 \times 32 \times 3$ pixels input's size) consisting of five layers, two convolutional layers including max pooling and three fully connected layers. The first and second convolutional layers consist of d_1 and d_2 filters, respectively, where $\frac{d_1}{d_2} \simeq \frac{6}{16}$. (B) The test error, ϵ , as a function of d_1 on a log-log scale, indicating a power-law scaling with exponent $\rho \sim 0.41$, Eq. (3) (Supplementary Information). The activation function of the nodes is ReLU.

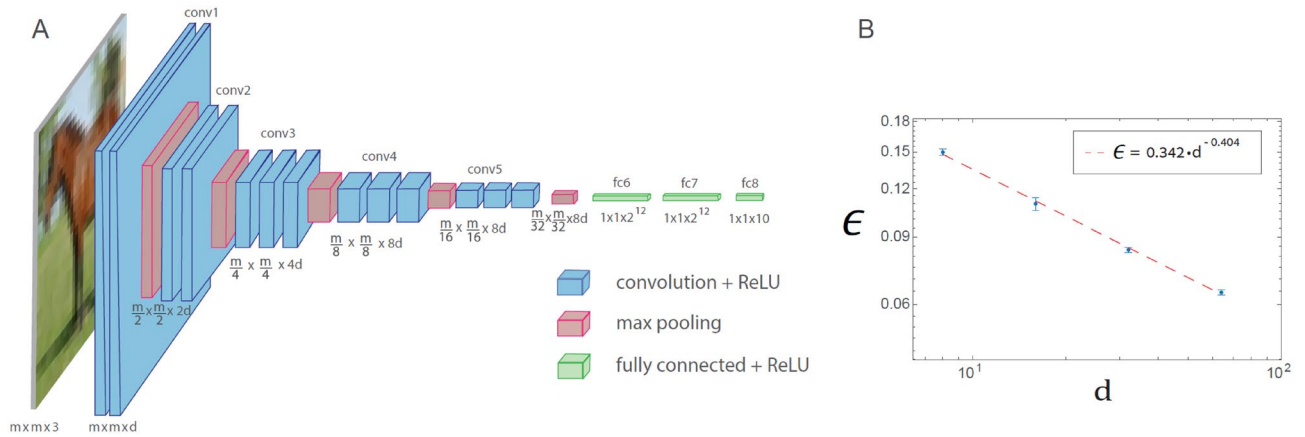


Figure 2. Learning in generalized VGG-16 architecture. **(A)** Generalized VGG-16 architecture consisting of 16 layers, where the number of filters in the n th convolution set is $d \cdot 2^{n-1}$ ($n \leq 4$) and the square-root of the size of the filter is $m \cdot 2^{-(n-1)}$ ($n \leq 5$), where $m \times m \times 3$ is the size of each input ($d = 64$ in the original VGG-16 architecture). **(B)** The test error, ϵ , as a function of d on a log–log scale, for the CIFAR-10 database ($m = 32$), indicating a power-law scaling with exponent $\rho \sim 0.4$, Eq. (3) (Supplementary Information). The activation function of the nodes is ReLU.

linearly with the number of filters. Hence, the number of weights increases linearly with the number of filters, whereas the number of weights in the successive fully connected layers remains constant and is independent of the number of filters (Figs. 1A, 2A).

The computational complexities as a function of the error rates were calculated using the power law extrapolation of $\epsilon(d)$ (Figs. 1B, 2B). The results indicate that the complexity increases with $1/\epsilon$, as a power law with an exponent ρ close to 5, ~ 4.85 for LeNet and ~ 4.94 for VGG-16 (Fig. 3B). Since error rates in both cases (Figs. 1B, 2B) are approximated by

$$\epsilon \propto \frac{1}{d^{0.4}},$$

therefore, $d \propto \epsilon^{-2.5}$ and in the leading term (Fig. 3A)

$$\text{Complexity} \propto d^2 \propto \epsilon^{-5}.$$

A direct calculation of the computational complexity ratio per step between LeNet and VGG-16, based on $\epsilon(d)$ (Figs. 1B, 2B), indicates that it is less than 0.6 for at least $\epsilon \geq 0.005$ (Fig. 3C). As it is extremely sensitive to the similar estimated values of ρ for both LeNet and VGG-16 (Figs. 1B, 2B), further extrapolation toward vanishing ϵ is unclear. Nevertheless, the lower complexity per epoch of shallow architectures serves as an example of the potential advantages of brain-inspired architectures. We note that the entire computational learning complexity is proportional to the number of training epochs and the classification of an input depends on a forward step only.

Under parallel computation, the required number of clock steps in a feedforward or BP realization is bounded from below by the number of layers. Decreasing this lower bound using a mechanism similar to that of carry-lookahead²⁷, developed for the addition and multiplication of large numbers, is practically inapplicable for such complex architectures. This is another expected advantage of learning based on brain-inspired shallow architecture.

The power-law behavior (Eq. 3) is demonstrated to govern both shallow and DL architectures, where the number of filters obeys the conservation law (Eq. 2). The following two questions were examined: The first question concerns the robustness of the power law (Eq. 3) for architectures that deviate from the conservation law (Eq. 2). The second question is whether Eq. (2), which controls the number of filters in the convolutional layers is indeed the optimized choice to minimize ϵ .

The power-law scaling for LeNet, which deviates from the conservation law (Eq. 2) is defined as follows:

$$\frac{d_2}{d_1} = \text{constant}, \tag{4}$$

which differs from $\frac{16}{6}$. For a smaller constant, $\frac{4}{3}$, the error rates were increased by a larger pre-factor A , as shown in Eq. (3); however, ρ remained similar ~ 0.4 (Fig. 4A). For a larger constant, $\frac{16}{3}$, the slope decreased, $\rho \sim 0.35$ (Fig. 4A, Supplementary Information). The results first indicate the robustness of the power law for various constants (Eq. 4) which alludes to its universal behavior. Second, for a smaller constant and any given d_1 , the error rates were enhanced. For a large constant (Eq. 4) and sufficiently large d_1 error rates were also enhanced because ρ decreased, but for small d_1 values, the error rates decreased. The results indicate that the conservation law (Eq. 2) with a constant that is expected to be approximately $\frac{16}{6}$, asymptotically minimizes ϵ for a large d_1 . Similar trends were obtained for VGG-16, where the number of filters in the n th convolution set

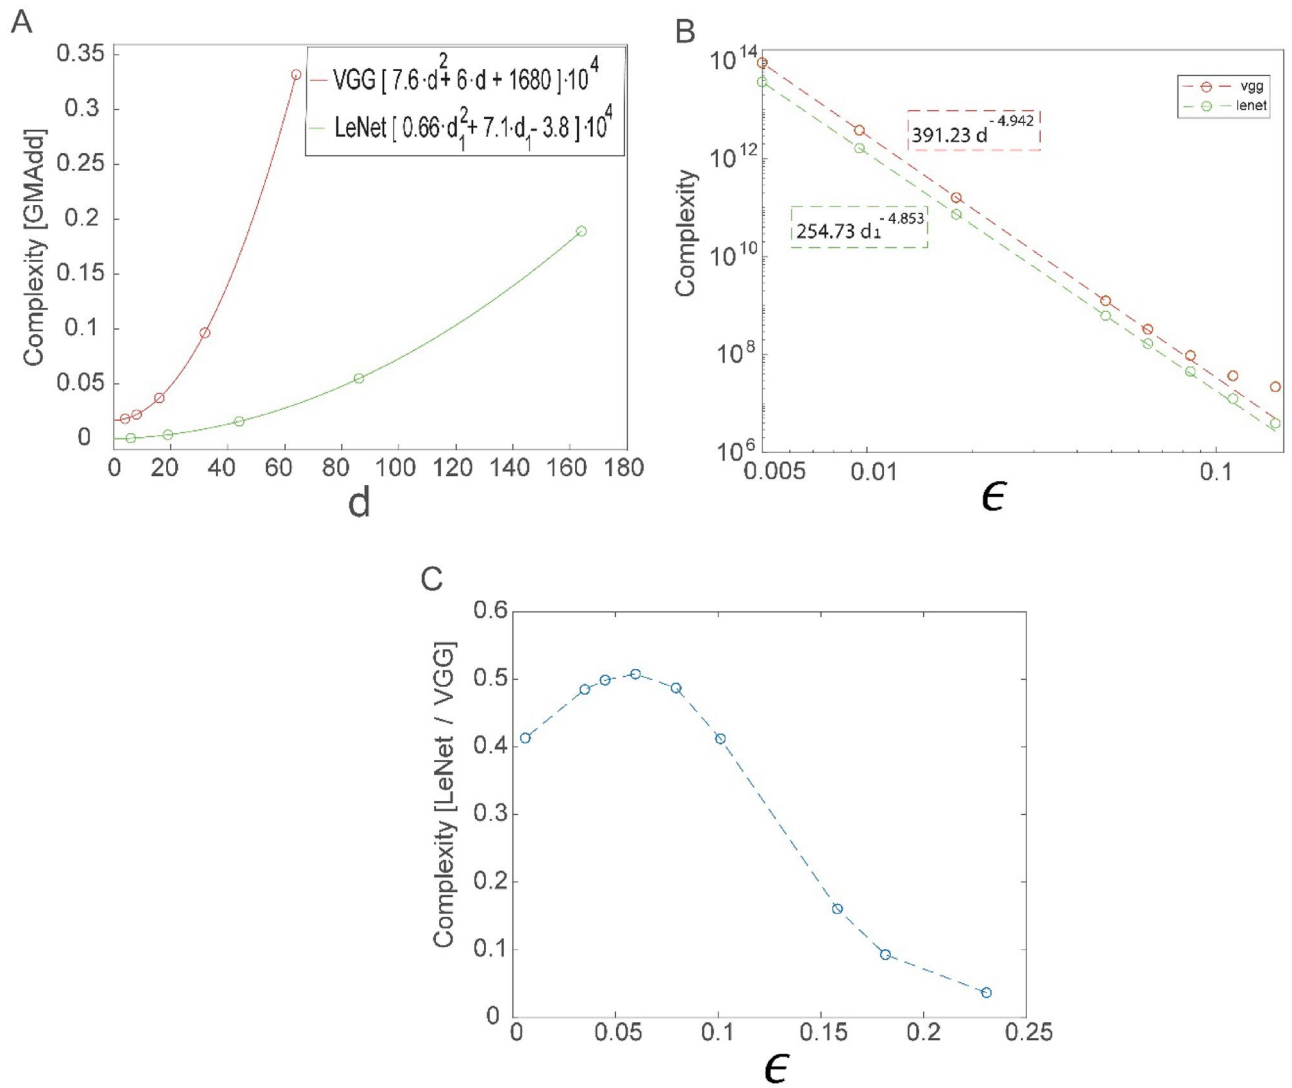


Figure 3. A comparison of learning complexity between generalized LeNet and VGG-16 architectures. (A) Complexity of a feedforward and BP step for a single input of LeNet (green) and VGG-16 (red) measured for several $d(d_1)$ values (open circles) and the quadratic polynomial fits. Complexity is measured in Giga multiplication-add operations (GMAdd). (B) Complexity as a function of $\epsilon(d)$ for LeNet (green) and VGG-16 (red) for several values of small ϵ (open circles) obtained from the power-law scaling (Figs. 1B, 2B) and the fitted power-law scaling (in dashed boxes), obtained from the last three small values of ϵ (dashed lines). (C) The ratio between the complexity of LeNet and VGG-16 for several values ϵ (open circles connected by a dashed line), obtained from the extrapolated $\epsilon(d_1)$ and $\epsilon(d)$ for LeNet and VGG-16 (Figs. 1B, 2B), respectively, and a direct measure of the complexity^{25,26} (Supplementary Information).

($n \leq 4$) increased as $constant^{(n-1)}$, whereas in the original architecture $constant = 2$ (Fig. 2A). For $constant = 1.5$, the error rates increased with a larger pre-factor, A , where ρ remained similar ~ 0.4 (Fig. 4B, Supplementary Information). For $constant = 2.5$, $\rho \sim 0.32$, indicating once more that the error rates increased asymptotically compared to $constant = 2$, but for small d , the error rates could be decreased. The results for VGG-16 indicate the robustness of the universal power-law behavior for various constants, as shown in Eq. (4), where a $constant$ close to 2 minimizes ϵ .

The following theoretical justification may explain why the conservation law (Eq. 2) leads to the minimization of error rates: Its purpose is to preserve the signal-to-noise ratio (SNR) along the feedforward convolutional layers such that the signal is repeatedly amplified. The noise of each large convolutional filter is expected to be proportional to the square-root of its size, m , and its signal to m^2 . Consequently, the SNR is proportional to m , and for the entire convolutional layer composed of depth d is $m \cdot d$. Hence, to compensate for the shrinking of the convolutional layer size along the feedforward architecture, its depth must be increased accordingly. Indeed, preliminary results indicate that doubling the number of filters in the fifth convolution set of VGG-16 (with $d = 16$), such that the number of filters in all convolutions ($n \leq 5$) is $16 \cdot 2^{n-1}$ decreased ϵ by ~ 0.015 compared with the standard VGG-16 architecture (Fig. 2, Supplementary Information for enhanced VGG-16). This supports the argument that maintaining the same SNR along the entire deep architecture enhances success

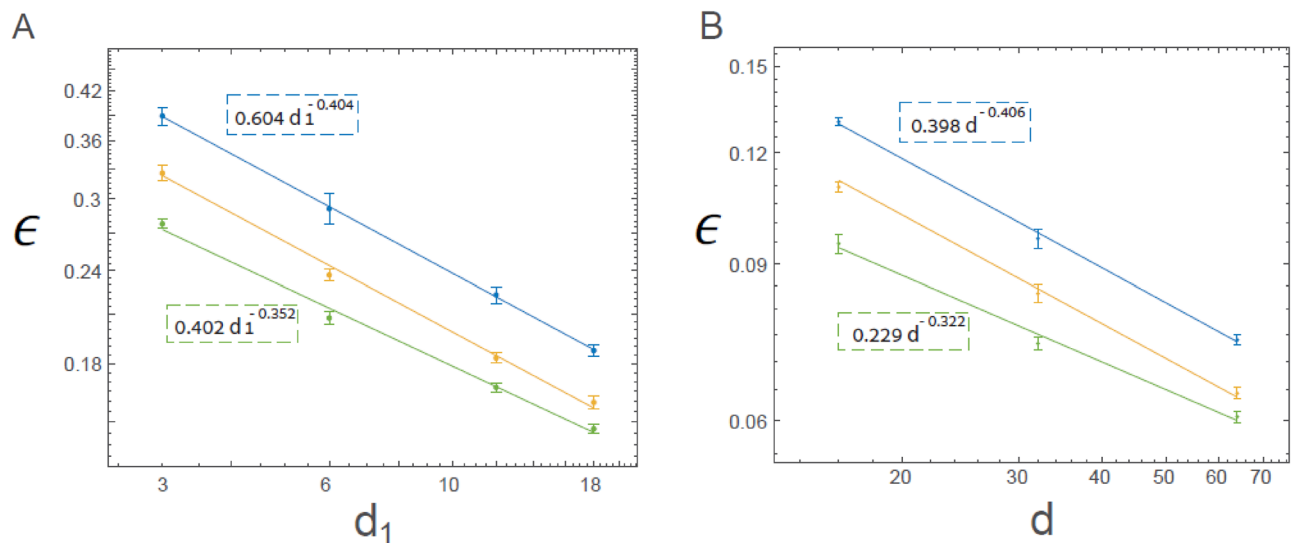


Figure 4. Conservation law indicating the optimal ratio between the depth of filters and their convolutional layer size. **(A)** Success rates and their standard deviation as a function of d_1 for the generalized LeNet architecture where $\frac{d_2}{d_1} = \frac{16}{3}$ (green) and $\frac{d_2}{d_1} = \frac{4}{3}$ (blue) with fitted (color coded) power law (Supplementary Information), and the results for $\frac{d_2}{d_1} = \frac{16}{6}$ (yellow, as shown in Fig. 1B) are presented for reference. **(B)** Similar to **(A)**, for the generalized VGG-16 architecture, where the number of filters in the n th convolution set ($n \leq 4$) is $d \cdot 2.5^{n-1}$ (green) and $d \cdot 1.5^{n-1}$ (blue) with fitted (color coded) power law (Supplementary Information), and results for $d \cdot 2^{n-1}$ (yellow, as shown in Fig. 2B) are presented for a reference.

rates. Nevertheless, further extended simulations on various architectures and databases are required to support the accuracy of the suggested conservation law, particularly because the convolutional layer sizes are small and far from the thermodynamic limit. Additionally, it is important to examine how the sensitivity of ρ and the conservation law are related to the properties of the cost function and the details of BP dynamics.

Discussion

Minimizing error rates for a particular classification task and database has been one of the primary goals of machine learning research over the past few decades. As a result, more structured DL architectures consisting of various combinations of concatenated convolutional and densely connected layers have been developed. Typically, further significant minimization of error rates requires deeper architectures, where such an architecture with some modifications can achieve reasonably high success rates for several other databases and classification tasks. This study suggests that, using the extrapolation of the power-law scaling (Eq. 3) traditional shallow architectures can achieve the same error rates as state-of-the-art DL architectures. The preferred architecture can reduce the space–time complexity for a specific training algorithm on a given database and hardware implementation. A theoretical framework is presented for constructing a hierarchy of complexity between families of artificial feedforward neural network architectures, based on their power-law scaling, exponent ρ , and pre-factor A . It is possible that the optimal architecture among several ones depends on the desired error rate (Fig. 4). Contrary to common knowledge, shallow feedforward brain-inspired architectures are not inferior, and they do not represent, as thought, an additional biological limitation²⁸. They can achieve low error rates such as DL algorithms, even with significantly low computational complexity for complex classification tasks (Fig. 3). We note that the presented power law as a function of the depth of the architecture differs from the power law behavior for SRs as a function of the dataset size^{29–33}.

Architectures that maximize ρ and its upper bound are not yet known. Preliminary results indicate that for a specific architecture, ρ may increase when the number of weights grows super-linearly with the number of filters. This can be achieved using a fully connected layer, in which the number of input and output units is proportional to the number of filters. Another possible mechanism is the addition of a super-linear number of cross-weights to the filters. This represents a biological realization because cross-weights result as a byproduct of dendritic nonlinear amplification^{17,29,34,35}. Nevertheless, these possible enhanced ρ mechanisms significantly increase computational complexity and are mentioned for their potential biological relevance, limited number of layers, and the natural emergence of many cross-weights.

Advanced GPU technology is used to minimize the running time of the DL algorithms. Indeed, our single-epoch running time using the CIFAR-10 database and VGG-16 with $d = 4$ is only a factor ~ 1.5 compared with LeNet with $d_1 = 6$, where both cases have similar success rates. However, shallow architectures with the same error rates as advanced deep architectures require more filters per convolutional layer, and consequently, a significantly increased number of fully connected weights. Above a critical number of filters, depending on the GPU properties, an epoch's running time is significantly slowed down and can even increase by a few orders of magnitude. Similarly, the running time in our case of VGG-16 with $d = 400$ is ~ 60 times slower than that

with $d = 8$, and LeNet with $d_1 = 2304$ is ~ 900 times slower than that with $d_1 = 6$. Hence, efficient realization of competitive error rates of shallow architectures to advance DL architectures requires a shift in the properties of advanced GPU technology. Additionally, it is expected to achieve a significant reduction in computational complexity for a desired error rate and a specific database (Fig. 3).

The power law behavior is presented in this work only for CIFAR10 database and its universal behavior must be confirmed in further research on other datasets. We note that this mission is difficult for the MNIST and CIFAR100 datasets. For MNIST, SRs exceed 0.99 even with LeNet and its power law extrapolation towards unity including error bars is improbable. On the other hand, for CIFAR100 the reported VGG16 success rates³⁶ are around ~ 0.74 , hence, the power law observation requires higher success rates such that finite size effects are minimized. However, extending the initial depth above $d = 64$ is beyond our computational capabilities.

The observation of power law as a function of the filter's depth must also be generalized to other architectures beyond LeNet and VGG16. We note that this mission requires careful optimization of systems' hyper-parameter space independently for each initial depth, resulting in a demand for high computational power. In addition, results are presented only for stochastic gradient descent algorithm^{21,22} and the robustness of the power law behavior needs to be verified on other more advanced optimizers³⁷, as well as its possible extension to other deep learning tasks, i.e. segmentation^{38,39}.

Additionally, for large dimension image inputs shallow networks exhibit degradation in SRs and it is commonly accepted that deep architectures are required to enhance SRs. The exhibited results suggest that increasing depth of shallow network will enhance SRs also for large dimension image inputs where the computational complexity is proportional to input size, as for deep architectures. However, its verification requires further research.

Finally, the theoretical origin of the universal power-law scaling (Eq. 3), governing shallow and DL architectures, has not yet been discovered. The following theoretical framework may provide a starting point for investigating this general phenomenon. The teacher–student online scenario is one of the analytically solvable cases exemplifying power law behavior cases⁴⁰. In the prototypical realizable scenario, the teacher and student have the same feedforward architecture, for example, a binary or a soft committee machine^{41–43}, but different initial weights. The teacher supports the students with a random input–output relation, and the student updates its weights based on this information and its current set of weights. The generalization error (test error) decays as a power law with the number of input–output examples, which is normalized to the size of the input. This work differs from online learning because the size of the non-random training dataset is limited, and a training example is repeatedly presented as an input without an online scenario. However, assuming a power-law scaling (Eq. 3), an architecture with an infinite number of filters $d \rightarrow \infty$, exists such that the test error vanishes. This architecture is the teacher's counterpart and represents a learning rule in the online scenario. A student with fewer filters attempts to imitate the teacher and results in a generalization error, which is expected to decrease with an increasing number of filters. It is currently impossible to find an analytical solution for the shallow and deep architectures that are shown as a function of the number of filters. The question is whether a toy model, where a filter may be represented by a perceptron with a nonlinear output unit, can be solved analytically to show that the generalization error decays as a power law with the number of filters.

Data availability

Source data are provided in this study, including all data supporting the plots, along with other findings of this study.

Received: 14 November 2022; Accepted: 29 March 2023

Published online: 20 April 2023

References

1. Hebb, D. O. *The Organization of Behavior: A Neuropsychological Theory* (Psychology Press, 2005).
2. Payeur, A., Guerguiev, J., Zenke, F., Richards, B. A. & Naud, R. Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nat. Neurosci.* **24**, 1010–1019 (2021).
3. Rosenblatt, F. *The Perceptron, a Perceiving and Recognizing Automaton Project Para* (Cornell Aeronautical Laboratory, 1957).
4. Minsky, M. & Papert, S. *Perceptrons Cambridge* (MIT Press, 1969).
5. Huang, G., Liu, Z., Van Der Maaten, L. & Weinberger, K. Q. *Proc. IEEE Conference on Computer Vision and Pattern Recognition* 4700–4708.
6. Han, D., Kim, J. & Kim, J. *Proc. IEEE Conference on Computer Vision and Pattern Recognition* 5927–5935.
7. LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2324 (1998).
8. Richards, B. A. *et al.* A deep learning framework for neuroscience. *Nat. Neurosci.* **22**, 1761–1770 (2019).
9. Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **25**, 1 (2012).
10. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
11. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
12. Fukushima, K. Cognitron: A self-organizing multilayered neural network. *Biol. Cybern.* **20**, 121–136 (1975).
13. LeCun, Y. & Bengio, Y. Convolutional networks for images, speech, and time series. *Handb. Brain Theory Neural Netw.* **3361**, 1995 (1995).
14. Serre, T. *et al.* *A Theory of Object Recognition: Computations and Circuits in the Feedforward Path of the Ventral Stream in Primate Visual Cortex* (Massachusetts Institute of Technology, 2005).
15. Hubel, D. H. & Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* **160**, 106 (1962).
16. Agliari, E., Alemanno, F., Barra, A. & De Marzo, G. The emergence of a concept in shallow neural networks. *Neural Netw.* **148**, 232–253 (2022).
17. Hodassman, S., Vardi, R., Tugendhaft, Y., Goldental, A. & Kanter, I. Efficient dendritic learning as an alternative to synaptic plasticity hypothesis. *Sci. Rep. U.K.* **12**, 1–12 (2022).

18. Bottou, L. *et al.* *Proc. 12th IAPR International Conference on Pattern Recognition, Vol. 3-Conference C: Signal Processing (Cat. No. 94CH3440-5)* 77–82 (IEEE).
19. LeCun, Y., Haffner, P., Bottou, L. & Bengio, Y. *Shape, Contour and Grouping in Computer Vision* 319–345 (Springer, 1999).
20. Krizhevsky, A. & Hinton, G. *Learning Multiple Layers of Features from Tiny Images* (2009).
21. Robbins, H. & Monro, S. A stochastic approximation method. *Ann. Math. Stat.* **22**, 400–407 (1951).
22. Kiefer, J. & Wolfowitz, J. Stochastic estimation of the maximum of a regression function. *Ann. Math. Stat.* **23**, 462–466 (1952).
23. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. Preprint at <http://arXiv.org/1409.1556> (2014).
24. Hasanpour, S. H., Rouhani, M., Fayyaz, M., Sabokrou, M. & Adeli, E. Towards principled design of deep convolutional networks: Introducing simpnet. Preprint at <http://arXiv.org/1802.06205> (2018).
25. Tschannen, M., Khanna, A. & Anandkumar, A. *International Conference on Machine Learning* 4985–4994 (PMLR).
26. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A. & Chen, L.-C. *Proc. IEEE Conference on Computer Vision and Pattern Recognition* 4510–4520.
27. Balasubramanian, P. & Mastorakis, N. Performance comparison of carry-lookahead and carry-select adders based on accurate and approximate additions. *Electronics* **7**, 369 (2018).
28. Hodassman, S. *et al.* Brain inspired neuronal silencing mechanism to enable reliable sequence identification. *Sci. Rep.* **12**, 16003. <https://doi.org/10.1038/s41598-022-20337-x> (2022).
29. Meir, Y. *et al.* Power-law scaling to assist with key challenges in artificial intelligence. *Sci. Rep. U.K.* **10**, 1–7 (2020).
30. Hestness, J. *et al.* Deep learning scaling is predictable, empirically. Preprint at <http://arXiv.org/1712.00409> (2017).
31. Sorscher, B., Geirhos, R., Shekhar, S., Ganguli, S. & Morcos, A. S. Beyond neural scaling laws: beating power law scaling via data pruning. Preprint at <http://arXiv.org/2206.14486> (2022).
32. Kaplan, J. *et al.* Scaling laws for neural language models. Preprint at <http://arXiv.org/2001.08361> (2020).
33. Henighan, T. *et al.* Scaling laws for autoregressive generative modeling. Preprint at <http://arXiv.org/2010.14701> (2020).
34. Sardi, S. *et al.* Adaptive nodes enrich nonlinear cooperative learning beyond traditional adaptation by links. *Sci. Rep.* **8**, 1–10 (2018).
35. Sardi, S. *et al.* Brain experiments imply adaptation mechanisms which outperform common AI learning algorithms. *Sci. Rep.* **10**, 1–10 (2020).
36. Chen, J., Sinha, S. & Kyriallidis, A. *Uncertainty in Artificial Intelligence* 326–335 (PMLR).
37. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. Preprint at <http://arXiv.org/1412.6980> (2014).
38. He, K., Zhang, X., Ren, S. & Sun, J. *Proc. IEEE Conference on Computer Vision and Pattern Recognition* 770–778.
39. Badrinarayanan, V., Kendall, A. & Cipolla, R. Segnet: A deep convolutional encoder–decoder architecture for image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**, 2481–2495 (2017).
40. Watkin, T. L., Rau, A. & Biehl, M. The statistical mechanics of learning a rule. *Rev. Mod. Phys.* **65**, 499 (1993).
41. Urbanczik, R. A fully connected committee machine learning unrealizable rules. *J. Phys. A Math. Gen.* **28**, 7097 (1995).
42. Schwarze, H., Oppen, M. & Kinzel, W. Generalization in a two-layer neural network. *Phys. Rev. A* **46**, R6185 (1992).
43. Biehl, M. & Schwarze, H. Learning by on-line gradient descent. *J. Phys. A: Math. Gen.* **28**, 643 (1995).

Acknowledgements

I.K. acknowledges the partial financial support from the Israel Science Foundation (Grant Number 346/22). S.H. acknowledges support from the Israeli Ministry of Science and Technology. We thank Amir Goldental and Roni Vardi for their stimulating discussions.

Author contributions

Y.M. contributed considerably to all simulations on LeNet and VGG-16, O.T. derived the main results for LeNet, R.G. contributed considerably to LeNet simulations, Y.T. contributed considerably to all simulations, and S.H. contributed to simulations and the preparation of the figures. I.K. initiated the study and supervised all aspects of the study. All authors have commented on the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-023-32559-8>.

Correspondence and requests for materials should be addressed to I.K.

Reprints and permissions information is available at 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/>.

© The Author(s) 2023