

Envisioning future deep learning theories: some basic concepts and characteristics

Weijie J. SU^{1,2}¹*The Wharton School, University of Pennsylvania, Philadelphia PA 19104, USA;*²*School of Engineering and Applied Science, University of Pennsylvania, Philadelphia PA 19104, USA*

Received 30 November 2023/Revised 8 March 2024/Accepted 5 August 2024/Published online 13 September 2024

Abstract To advance deep learning methodologies in the next decade, a theoretical framework for reasoning about modern neural networks is needed. While efforts are increasing toward demystifying why deep learning is so effective, a comprehensive picture remains lacking, suggesting that a better theory is possible. We argue that a future deep learning theory should inherit three characteristics: a hierarchically structured network architecture, parameters iteratively optimized using stochastic gradient-based methods, and information from the data that evolves compressively. As an instantiation, we integrate these characteristics into a graphical model called neurashed. This model effectively explains some common empirical patterns in deep learning. In particular, neurashed enables insights into implicit regularization, information bottleneck, and local elasticity. Finally, we discuss how neurashed can guide the development of deep learning theories.

Keywords deep learning, local elasticity, information bottleneck, implicit regularization, phenomenological model

1 Introduction

Deep learning is recognized as a monumentally successful approach to many data-extensive applications in image recognition, natural language processing, and board game programs [1–3]. Despite extensive efforts [4–6], however, our theoretical understanding of how this increasingly popular machinery works and why it is so effective remains incomplete. This is exemplified by the substantial vacuum between the highly sophisticated training paradigm of modern neural networks and the capabilities of existing theories. For instance, the optimal architectures for certain specific tasks in computer vision remain unclear [7].

To better fulfill the potential of deep learning methodologies in increasingly diverse domains, heuristics and computation are unlikely to be adequate—a comprehensive theoretical foundation for deep learning is needed. Ideally, this theory would demystify these black-box models, visualize the essential elements, and enable principled model design and training. A useful theory would, at a minimum, reduce unnecessary computational burden and human costs in present-day deep-learning research, even if it could not make all complex training details transparent.

Unfortunately, it is unclear how to develop a deep learning theory from first principles. Instead, in this paper we take a phenomenological approach that captures some important characteristics of deep learning. Roughly speaking, a phenomenological model provides an overall picture rather than focusing on details, and allows for useful intuition and guidelines so that a more complete theoretical foundation can be developed.

To address what characteristics of deep learning should be considered in a phenomenological model, we recall the three key components in deep learning: architecture, algorithm, and data [8]. The most pronounced characteristic of modern network architectures is their hierarchical composition of simple functions. Indeed, overwhelming evidence shows that multiple-layer architectures are superior to their shallow counterparts [9], reflecting the fact that high-level features are hierarchically represented through




Figure 1 Neurashed model that imitates a four-layer neural network for a three-class classification problem. For instance, the feature represented by the leftmost node in the second level is formed by the features represented by the three leftmost nodes in the first level.

low-level features [10,11]. The optimization workhorse for training neural networks is stochastic gradient descent or Adam [12], which iteratively updates the network weights using noisy gradients evaluated from small batches of training samples. Overwhelming evidence shows that the solution trajectories of iterative optimization are crucial to generalization performance [13]. It is also known that the effectiveness of deep learning relies heavily on the structure of the data [14,15], which enables the compression of data information in the late stages of deep learning training [16,17].

2 Neurashed

We introduce a simple, interpretable, white-box model that simultaneously possesses the hierarchical, iterative, and compressive characteristics to guide the development of a future deep learning theory. This model, called neurashed, is a phenomenological model that mimics neural networks through the training process. Depending on the neural network that it imitates, a neurashed model is represented as a graph with nodes partitioned into different levels (Figure 1). For example, the number of levels is the same as the number of layers of the corresponding neural network. Instead of corresponding with a single neuron in the corresponding neural network, an l -level node in neurashed represents a feature that the neural network can learn in its l -layer. For example, the nodes in the first/bottom level denote lowest-level features, whereas the nodes in the last/top level correspond to the class membership in the classification problem. To describe the dependence of high-level features on low-level features, neurashed includes edges between a node and its dependent nodes in the preceding level. This reflects the hierarchical nature of features in neural networks.

Given any input sample, a node in neurashed is in one of two states: firing or not firing. The first-layer nodes take the pixels as input when the input is an image¹⁾. The unique last-level node that fires for an input corresponds with the label of the input. Whether a node in the first level fires or not is determined by the input. For a middle-level node, its state is determined by the firing pattern of its dependent nodes in the preceding levels through its firing rule. For example, let a node represent **cat** and its dependent nodes be **cat head** and **cat tail**. We can activate **cat** when either or both of the two dependent nodes are firing. Alternatively, let a node represent **panda head** and consider its dependent nodes **dark circle**, **black ear**, and **white face**. Let the **panda head** node be firing only if all three dependent nodes are firing. In general, the firing rule is a Boolean function with arity being the number of dependent nodes.

We call the subgraph induced by the firing nodes the feature pathway of a given input. A feature pathway should be consistent with the firing rules. In general, samples from different classes have relatively distinctive feature pathways, commonly shared at lower levels but more distinct at higher levels. By contrast, feature pathways of same-class samples are identical or similar. An illustration is given in Figure 2 [18].

To mimic the prediction of neural networks, all nodes F in neurashed except for the last-level nodes are assigned a nonnegative value λ_F as a measure of each node’s ability to sense the corresponding feature. A large value of λ_F means that when this node fires it can send out strong signals to connected nodes in the next level. Hence, λ_F is the amplification factor of F . Moreover, let η_{fF} denote the weight of a connected second-last-level node f and last-level node F . Given an input, we define the score of each node, which is

1) For simplicity, we assume that a node is firing if the corresponding pixel value is above a certain threshold.




Figure 2 (Color online) Feature pathways of the neurashed model in Figure 1. Firing nodes are marked in red. Class 1 includes two types of samples with slightly different feature pathways, which is a reflection of heterogeneity in real-life data [18]. (a) Class 1a; (b) Class 1b; (c) Class 2; (d) Class 3.

sent to its connected nodes on the next level: For any first-level node F , let score $S_F = \lambda_F$ if F is firing and $S_F = 0$ otherwise; for any firing middle-level node F , its score depends on its amplification factor as well as the scores of its dependent nodes:

$$S_F = \lambda_F \sum_{f \rightarrow F} S_f,$$

where the sum is over all dependent nodes f of F in the lower level. Likewise, let $S_F = 0$ for any non-firing middle-level node F . For the last-level nodes F_1, \dots, F_K corresponding to the K classes, let

$$Z_j = \sum_{f \rightarrow F_j} \eta_{fF_j} S_f \quad (1)$$

be the logit for the j th class, where the sum is over all second-last-level dependent nodes f of F_j . Finally, we predict the probability that this input is in the j th class as

$$p_j(x) = \frac{\exp(Z_j)}{\sum_{i=1}^K \exp(Z_i)}.$$

To mimic the iterative characteristic of neural network training, we must be able to update the amplification factors for neurashed in correspondence with backpropagation in the neural network. At initialization, because there is no predictive ability as such for neurashed, we set λ_F and η_{fF} to zero, other constants, or random numbers. We assume that a node is firing if it is in the union of the feature pathways of all training samples in the mini-batch for computing the gradient. In each backpropagation, the neural network can improve its ability to sense features of images in the mini-batch. In light of this, we increase the amplification ability of any firing node, thereby facilitating the predictions for samples with similar feature pathways. This recognizes the famous adage “cells that fire together wire together” from Hebbian learning [19]. Specifically, if a node F is firing in the backpropagation, we update its amplification factor λ_F by letting

$$\lambda_F \leftarrow g^+(\lambda_F), \quad (2)$$

where g^+ is an increasing function satisfying $g^+(x) > x$ for all $x \geq 0$. The simplest choices include $g^+(x) = ax$ for $a > 1$ and $g^+(x) = x + c$ for $c > 0$. The strengthening of firing feature pathways is consistent with a recent analysis of simple hierarchical models [20, 21]. By contrast, for any node F that is not firing in the backpropagation, we decrease its amplification factor by setting

$$\lambda_F \leftarrow g^-(\lambda_F) \quad (3)$$

Algorithm 1 Training dynamics of neurashed

Input: neurashed graph, update rule g^+ and g^- , input data with feature pathways;
while training continues **do**
 Sample a mini-batch of size B ;
 Let G be the union of all feature pathways of the mini-batch sample;
 for each node F in neurashed **do**
 if $F \in G$ **then**
 $\lambda_F \leftarrow g^+(\lambda_F)$;
 else
 $\lambda_F \leftarrow g^-(\lambda_F)$;
 end if
 end for
end while
Output: amplification factors of all nodes.

for an increasing function g^- satisfying $0 \leq g^-(x) \leq x$; for example, $g^-(x) = bx$ for some $0 < b \leq 1$. This recognizes regularization techniques such as weight decay, batch normalization [22], layer normalization [23], and dropout [24] in deep-learning training, which effectively impose certain constraints on the weight parameters [25]. Update rules g^+, g^- generally vary with respect to nodes and iteration number. Likewise, we apply rule g^+ to η_{fF} when the connected second-last-level node f and last-level node F both fire; otherwise, g^- is applied. An illustration is given in Algorithm 1.

The training dynamics above could improve neurashed’s predictive ability. In particular, the update rules allow nodes appearing frequently in feature pathways to quickly grow their amplification factors. Consequently, updating an image would improve the prediction on images with similar feature pathways, which is consistent with the local elasticity phenomenon in real-world neural networks [26]. Moreover, for an input x belonging to the j th class, the amplification factors of most nodes in its feature become relatively large during training, and the true-class logit Z_j also becomes much larger than the other logits Z_i for $i \neq j$. This shows that the probability of predicting the correct class $p_j(x) \rightarrow 1$ as the number of iterations tends to infinity.

While neurashed may appear to have all features predetermined before training and not capture the extraction of features, particularly for nodes in the bottom layer, this limitation is less critical than it might seem. In the neurashed model, one can envision the graph as containing a vast array of potential feature pathways, with the training dynamics described by the update rules (2) and (3) serving to identify the relevant features among a plethora of potentially irrelevant ones. This perspective is consistent with the random feature view of neural networks [27, 28], in which features remain fixed during training, and only the feature weights are updated. Meanwhile, it is worth mentioning that a crucial distinction between our neurashed and the random feature model is that the former allows for hierarchy in feature formation.

The modeling strategy of neurashed is similar to a watershed, where tributaries meet to form a larger stream (hence “neurashed”). This modeling strategy inherently endows neurashed with a hierarchical structure and facilitates iterative optimization. As a caveat, we do not regard the feature representation of neurashed as fixed. Although the graph is fixed, the evolving amplification factors represent features in a dynamic manner. Note that neurashed is different from capsule networks [29] and GLOM [10] in that our model is meant to shed light on the black box of deep learning, not serve as a working system.

3 Insights into puzzles

Implicit regularization. Conventional wisdom from statistical learning theory suggests that a model may not perform well on test data if its parameters outnumber the training samples; to avoid overfitting, explicit regularization is needed to constrain the search space of the unknown parameters [30]. In contrast to other machine learning approaches, modern neural networks—where the number of learnable parameters is often orders of magnitude larger than that of the training samples—enjoy surprisingly good generalization even without explicit regularization [31]. From an optimization viewpoint, this shows that simple stochastic gradient-based optimization for training neural networks implicitly induces a form of regularization biased toward local minima of low “complexity” [13, 32]. However, it remains unclear how implicit regularization occurs from a geometric perspective [33–35].

To gain geometric insights into implicit regularization using our conceptual model, recall that only firing features grow during neurashed training, whereas the remaining features become weaker during




Figure 3 (Color online) Part of neurashed that corresponds to a single class. Feature pathways of (a) Class 1 and (b) Class 2 using small-batch training (the last-level (layer three) node is firing but is not marked in red for simplicity); (c) the learned neurashed models, where larger nodes indicate larger amplification factors and the three nodes in the middle indicate a sparse learned feature pathway.




Figure 4 (Color online) (a) Firing pattern when both feature pathways are included in the case of large-batch training (as in Figure 3, the last-level node is not marked in red for simplicity). (b) A dense feature pathway is learned using large-batch training compared to Figure 3.

backpropagation. For simplicity, consider stochastic gradient descent with a mini-batch size of 1. Here, only common features shared by samples from different classes constantly fire in neurashed, whereas features peculiar to some samples or certain classes fire less frequently. As a consequence, these common features become stronger more quickly, whereas the other features grow less rapidly or even diminish.

When gradient descent or large-batch stochastic gradient descent is used, many features fire in each update of neurashed, thereby increasing their amplification factors simultaneously. By contrast, a small-batch method constructs the feature pathways in a sparing way. Consequently, the feature pathways learned using small batches are sparser, suggesting a form of compression. This comparison is illustrated in Figures 3 and 4, which imply that different samples from the same class tend to exhibit vanishing variability in their high-level features during later training, and is consistent with the recently observed phenomenon of neural collapse [36]. Intuitively, this connection is indicative of neurashed’s compressive nature. The law of data separation further extends this connection across layers [37].

Although the concept that small-batch training can lead to implicit regularization remains a hypothesis, there is considerable supporting evidence, both empirical and theoretical. Empirical studies in [38, 39] showed that neural networks trained by small-batch methods generalize better than when trained by large-batch methods. Moreover, Ref. [40, 41] showed that neural networks tend to be more accurate on test data if these models leverage less information of the images. From a theoretical angle, Ref. [42] related generalization performance to a solution’s sparsity level when a simple nonlinear model is trained using stochastic gradient descent.

Information bottleneck. In [16, 17], the information bottleneck theory of deep learning was introduced, based on the observation that neural networks undergo an initial fitting phase followed by a compression phase. In the initial phase, neural networks seek to both memorize the input data and fit the labels, as manifested by the increase in mutual information—a measure quantifying the reduction in uncertainty about one random variable by observing another—between a hidden level and both the

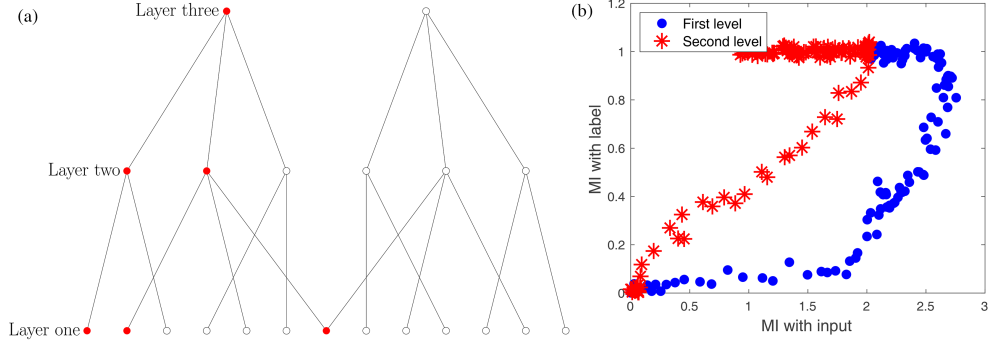


Figure 5 (Color online) Neurashed model for a binary classification problem. In the current plot, the 1st, 2nd, and 7th nodes (from left to right) on the first level are firing, which are denoted as (1, 2, 7). In total, there are four firing patterns of Class 1 on the first level: (1, 2, 7), (2, 3, 7), (4, 5, 7), (5, 6, 7). In the second level, the first and third nodes fire if one or more dependent nodes fire, and the second (dominant) node fires if two or more dependent nodes fire. (a) A feature pathway of Class 1. Class 2 has four feature pathways that are symmetric to those of Class 1. (b) The information bottleneck phenomenon for this neurashed model. As with [17], noise is added in calculating the mutual information (MI) between the first/second level and the input (8 types)/labels (2 types). More details are given in Appendix A.

input and labels. In the second phase, the networks compress all irrelevant information from the input, as demonstrated by the decrease in mutual information between the hidden level and input.

Instead of explaining how this mysterious phenomenon emerges in deep learning, which is beyond our scope, we shed some light on the information bottleneck by producing the same phenomenon using neurashed. As with implicit regularization, we observe that neurashed usually contains many redundant feature pathways when learning class labels. Initially, many nodes grow and thus encode more information regarding both the input and class labels. Subsequently, more frequently firing nodes become more dominant than less frequently firing ones. Because nodes compete to grow their amplification factors, dominant nodes tend to dwarf their weaker counterparts after a sufficient amount of training. Hence, neurashed starts to “forget” the information encoded by the weaker nodes, thereby sharing less mutual information with the input samples (see an illustration in Figure 5). The compressive characteristic of neurashed arises, loosely speaking, from the internal competition among nodes. This interpretation of the information bottleneck via neurashed is reminiscent of the human brain, which has many neuron synapses during childhood that are pruned to leave fewer firing connections in adulthood [43].

Local elasticity. Last, we consider a recently observed phenomenon termed local elasticity [26] in deep learning training, which asks how the update of neural networks via backpropagation at a base input changes the prediction at a test sample. Formally, for K -class classification, let $z_1(x, w), \dots, z_K(x, w)$ be the logits prior to the softmax operation with input x and network weights w . Writing w^+ for the updated weights using the base input x , we define

$$\text{LE}(x, x') := \frac{\sqrt{\sum_{i=1}^K (z_i(x', w^+) - z_i(x', w))^2}}{\sqrt{\sum_{i=1}^K (z_i(x, w^+) - z_i(x, w))^2}} \quad (4)$$

as a measure of the impact of base x on test x' . A large value of this measure indicates that the base has a significant impact on the test input. Through extensive experiments, Ref. [26] demonstrated that well-trained neural networks are locally elastic in the sense that the value of this measure depends on the semantic similarity between two samples x and x' . If they are similar—say, images of a cat and tiger—the impact is significant, and if they are dissimilar—say, images of a cat and turtle—the impact is low. Experimental results are shown in Figure 6 [44, 45]. For comparison, local elasticity does not appear in linear classifiers because of the leverage effect, which refers to the disproportionate influence of an observation on the fitted value of a least-squares model when that observation lies far from the bulk of the explanatory variables. More recently, Refs. [46–48] showed that local elasticity implies good generalization ability.

We now show that neurashed exhibits the phenomenon of local elasticity, which yields insights into how local elasticity emerges in deep learning. To see this, note that similar training samples share more of their feature pathways. For example, the two types of samples in Class 1 in Figure 2 are presumably very similar and indeed have about the same feature pathways; Class 1 and Class 2 are more

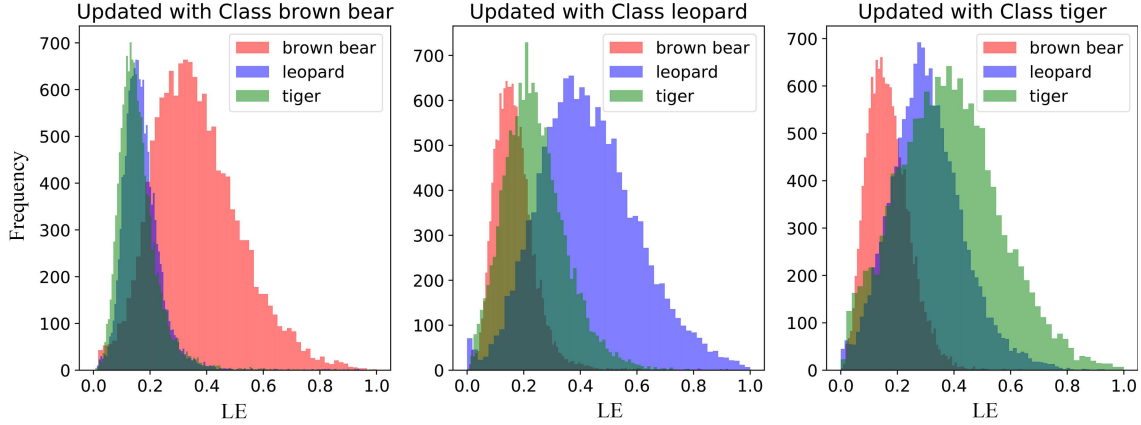


Figure 6 (Color online) Histograms of $LE(x, x')$ defined in (4), evaluated on the pre-trained VGG-19 network [44]. For example, in the left panel the base input x are images of brown bears. Each class contains 120 images sampled from ImageNet [45]. Tigers and leopards are both felines and are therefore similar. The histograms show that local elasticity appears to be larger when the two classes are the same or similar.

similar to each other than Class 1 and Class 3 in terms of feature pathways. Metaphorically speaking, applying backpropagation to an image of a leopard, the feature pathway for the leopard strengthens as the associated amplification factors increase. While this update also strengthens the feature pathway for the tiger, it does not impact the brown bear feature pathway as much, which presumably overlaps less with the leopard feature pathway. This update in turn leads to a more significant change in the logits (1) of an image of a tiger than those of a brown bear. Returning to Figure 2 for an illustration of this interpretation, the impact of updating at a sample in Class 1a is most significant on Class 1b, less significant on Class 2, and unnoticeable on Class 3.

4 Outlook

In addition to shedding new light on implicit regularization, information bottleneck, and local elasticity, neurashed is likely to facilitate insights into other common empirical patterns of deep learning. First, a byproduct of our interpretation of implicit regularization might evidence a subnetwork with comparable performance to the original, which could have implications on the lottery ticket hypothesis of neural networks [49]. Second, while a significant fraction of classes in ImageNet [45] have fewer than 500 training samples, deep neural networks perform well on these classes in tests. Neurashed could offer a new perspective on these seemingly conflicting observations—many classes are basically the same (for example, ImageNet contains 120 dog-breed classes), so the effective sample size for learning the common features is much larger than the size of an individual class. Last, neurashed might help reveal the benefit of data augmentation techniques such as cropping. In the language of neurashed, `cat head` and `cat tail` each are sufficient to identify `cat`. If both concepts appear in the image, cropping reinforces the neurashed model by impelling it to learn these concepts separately. Nevertheless, these views are preliminary and require future consolidation.

While closely resembling neural networks in many aspects, neurashed is not merely intended to better explain some phenomena in deep learning. Instead, our main goal is to offer insights into the development of a comprehensive theoretical foundation for deep learning in future research. In particular, neurashed’s efficacy in interpreting many puzzles in deep learning could imply that neural networks and neurashed evolve similarly during training. We therefore believe that a comprehensive deep learning theory is unlikely without incorporating the hierarchical, iterative, and compressive characteristics. That said, useful insights can be derived from analyzing models without these characteristics in some specific settings [4, 50–58].

Integrating the three characteristics in a principled manner might necessitate a novel mathematical framework for reasoning about the composition of nonlinear functions. Because it could take years before such mathematical tools become available, a practical approach for the present, given that such theoretical guidelines are urgently needed [59], is to better relate neurashed to neural networks and develop finer-grained models. For example, an important question is to understand the correspondence between a

neural network and a neurashed model and, in particular, how they speak for each other. Specifically, we need to determine the unit in neural networks that corresponds with a feature node in neurashed. For example, a feature node might correspond to a few neurons in the same layer of the neural network that are often activated at the same time. Indeed, recent studies showed that activation maps are sparse in modern architectures in the sense that only a very small portion of neurons are activated in prediction [60, 61]. This can serve as a good starting point for investigating the relationship between neuron activations in neural networks and feature pathways in neurashed, and more specifically, to explore the possibility of using neurashed to improve the interpretability of neural networks in practice. To generalize neurashed, edges could be fired instead of nodes. Another potential extension is to introduce stochasticity to rules g^+ and g^- for updating amplification factors and rendering feature pathways random or adaptive to learned amplification factors. Relaxing the constraint that the connection of the neurashed graph is fixed could also be an interesting avenue to explore. Owing to the flexibility of neurashed as a graphical model, such possible extensions are endless.

Acknowledgements This work was supported in part by an Alfred Sloan Research Fellowship and the Wharton Dean's Research Fund. We would like to thank Patrick CHAO, Zhun DENG, Cong FANG, Hangfeng HE, Qingxuan JIANG, Konrad KORDING, Yi MA, and Jiayao ZHANG for their helpful discussions and comments. We are grateful to two anonymous reviewers for their constructive comments that helped improve the presentation of the paper.

References

- 1 Krizhevsky A, Sutskever I, Hinton G E. ImageNet classification with deep convolutional neural networks. *Commun ACM*, 2017, 60: 84–90
- 2 LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*, 2015, 521: 436–444
- 3 Silver D, Huang A, Maddison C J, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 2016, 529: 484–489
- 4 Jacot A, Gabriel F, Hongler C. Neural tangent kernel: convergence and generalization in neural networks. In: *Proceedings of Advances in Neural Information Processing Systems*, 2018. 8580–8589
- 5 Bartlett P L, Foster D J, Telgarsky M. Spectrally-normalized margin bounds for neural networks. In: *Proceedings of Advances in Neural Information Processing Systems*, 2017. 6241–6250
- 6 Berner J, Grohs P, Kutyniok G, et al. The modern mathematics of deep learning. 2021. ArXiv:2105.04026
- 7 Tolstikhin I, Housley N, Kolesnikov A, et al. MIP-Mixer: an all-MLP architecture for vision. 2021. ArXiv:2105.01601
- 8 Zdeborová L. Understanding deep learning is also a job for physicists. *Nat Phys*, 2020, 16: 602–604
- 9 Eldan R, Shamir O. The power of depth for feedforward neural networks. In: *Proceedings of Conference on Learning Theory*, 2016. 907–940
- 10 Hinton G. How to represent part-whole hierarchies in a neural network. 2021. ArXiv:2102.12627
- 11 Bagrov A A, Iakovlev I A, Iliashov A A, et al. Multiscale structural complexity of natural patterns. *Proc Natl Acad Sci USA*, 2020, 117: 30241–30251
- 12 Kingma D P, Ba J. Adam: a method for stochastic optimization. In: *Proceedings of International Conference on Learning Representations*, 2015
- 13 Soudry D, Hoffer E, Nacson M S, et al. The implicit bias of gradient descent on separable data. *J Mach Learn Res*, 2018, 19: 2822–2878
- 14 Blum A L, Rivest R L. Training a 3-node neural network is NP-complete. *Neural Networks*, 1992, 5: 117–127
- 15 Goldt S, Mézard M, Krzakala F, et al. Modeling the influence of data structure on learning in neural networks: the hidden manifold model. *Phys Rev X*, 2020, 10: 041044
- 16 Tishby N, Zaslavsky N. Deep learning and the information bottleneck principle. In: *Proceedings of IEEE Information Theory Workshop (ITW)*, 2015. 1–5
- 17 Shwartz-Ziv R, Tishby N. Opening the black box of deep neural networks via information. 2017. ArXiv:1703.00810
- 18 Feldman V. Does learning require memorization? A short tale about a long tail. In: *Proceedings of Symposium on Theory of Computing*, 2020. 954–959
- 19 Hebb D O. *The Organization of Behavior: A Neuropsychological Theory*. New York: Psychology Press, 2005
- 20 Poggio T, Banburski A, Liao Q. Theoretical issues in deep networks. *Proc Natl Acad Sci USA*, 2020, 117: 30039–30045
- 21 Allen-Zhu Z, Li Y. Backward feature correction: how deep learning performs deep learning. 2020. ArXiv:2001.04413
- 22 Ioffe S, Szegedy C. Batch normalization: accelerating deep network training by reducing internal covariate shift. In: *Proceedings of International Conference on Machine Learning*, 2015. 448–456
- 23 Ba J L, Kiros J R, Hinton G E. Layer normalization. 2016. ArXiv:1607.06450
- 24 Srivastava N, Hinton G, Krizhevsky A, et al. Dropout: a simple way to prevent neural networks from overfitting. *J Mach Learn Res*, 2014, 15: 1929–1958
- 25 Fang C, He H, Long Q, et al. Exploring deep neural networks via layer-peeled model: minority collapse in imbalanced training. In: *Proceedings of the National Academy of Sciences*, 2021
- 26 He H F, Su W J. The local elasticity of neural networks. In: *Proceedings of International Conference on Learning Representations*, 2020
- 27 Rahimi A, Recht B. Random features for large-scale kernel machines. In: *Proceedings of Advances in Neural Information Processing Systems*, 2007. 20
- 28 Yehudai G, Shamir O. On the power and limitations of random features for understanding neural networks. In: *Proceedings of Advances in Neural Information Processing Systems*, 2019. 32
- 29 Sabour S, Frosst N, Hinton G E. Dynamic routing between capsules. In: *Proceedings of Advances in Neural Information Processing Systems*, 2017. 3859–3869
- 30 Friedman J, Hastie T, Tibshirani R. *The Elements of Statistical Learning*. New York: Springer, 2001
- 31 Zhang C Y, Bengio S, Hardt M, et al. Understanding deep learning (still) requires rethinking generalization. *Commun ACM*, 2021, 64: 107–115
- 32 Bartlett P L, Long P M, Lugosi G, et al. Benign overfitting in linear regression. *Proc Natl Acad Sci USA*, 2020, 117: 30063–30070
- 33 Nagarajan V, Kolter J Z. Uniform convergence may be unable to explain generalization in deep learning. In: *Proceedings of Advances in Neural Information Processing Systems*, 2019
- 34 Razin N, Cohen N. Implicit regularization in deep learning may not be explainable by norms. In: *Proceedings of Advances in Neural Information Processing Systems*, 2020. 33

- 35 Zhou Z-H. Why over-parameterization of deep neural networks does not overfit? *Sci China Inf Sci*, 2021, 64: 116101
- 36 Pappayan V, Han X Y, Donoho D L. Prevalence of neural collapse during the terminal phase of deep learning training. *Proc Natl Acad Sci USA*, 2020, 117: 24652–24663
- 37 He H F, Su W J. A law of data separation in deep learning. *Proc Natl Acad Sci USA*, 2023, 120: e2221704120
- 38 Keskar N S, Mudigere D, Nocedal J, et al. On large-batch training for deep learning: generalization gap and sharp minima. 2016. ArXiv:1609.04836
- 39 Smith S, Elsen E, De S. On the generalization benefit of noise in stochastic gradient descent. In: *Proceedings of International Conference on Machine Learning*, 2020. 9058–9067
- 40 Ilyas A, Santurkar S, Engstrom L, et al. Adversarial examples are not bugs, they are features. In: *Proceedings of Advances in Neural Information Processing Systems*, 2019. 248–255
- 41 Xiao K, Engstrom L, Ilyas A, et al. Noise or signal: the role of image backgrounds in object recognition. In: *Proceedings of International Conference on Learning Representations*, 2021
- 42 HaoChen J Z, Wei C, Lee J D, et al. Shape matters: understanding the implicit bias of the noise covariance. 2020. ArXiv:2006.08680
- 43 Feinberg I. Schizophrenia: caused by a fault in programmed synaptic elimination during adolescence? *J Psychiatric Res*, 1982, 17: 319–334
- 44 Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. 2014. ArXiv:1409.1556
- 45 Deng J, Dong W, Socher R, et al. ImageNet: a large-scale hierarchical image database. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2009. 248–255
- 46 Chen S, He H, Su W J. Label-aware neural tangent kernel: toward better generalization and local elasticity. In: *Proceedings of Advances in Neural Information Processing Systems*, 2020. 15847–15858
- 47 Deng Z, He H F, Su W J. Toward better generalization bounds with locally elastic stability. In: *Proceedings of International Conference on Machine Learning*, 2021
- 48 Zhang J Y, Wang H, Su W J. Imitating deep learning dynamics via locally elastic stochastic differential equations. In: *Proceedings of Advances in Neural Information Processing Systems*, 2021
- 49 Frankle J, Carbin M. The lottery ticket hypothesis: finding sparse, trainable neural networks. In: *Proceedings of International Conference on Learning Representations*, 2018
- 50 Chizat L, Oyallon E, Bach F. On lazy training in differentiable programming. In: *Proceedings of Advances in Neural Information Processing Systems*, 2019
- 51 Wu L, Ma C, E W. How SGD selects the global minima in over-parameterized learning: a dynamical stability perspective. In: *Proceedings of Advances in Neural Information Processing Systems*, 2018. 8289–8298
- 52 Mei S, Montanari A, Nguyen P M. A mean field view of the landscape of two-layer neural networks. *Proc Natl Acad Sci USA*, 2018, 115: 7665–7671
- 53 Chizat L, Bach F. On the global convergence of gradient descent for over-parameterized models using optimal transport. In: *Proceedings of Advances in Neural Information Processing Systems*, 2018. 3040–3050
- 54 Belkin M, Hsu D, Ma S, et al. Reconciling modern machine-learning practice and the classical bias-variance trade-off. *Proc Natl Acad Sci USA*, 2019, 116: 15849–15854
- 55 Lee J, Xiao L, Schoenholz S S, et al. Wide neural networks of any depth evolve as linear models under gradient descent. 2019. ArXiv:1902.06720
- 56 Xu Z Q J, Zhang Y Y, Xiao Y Y. Training behavior of deep neural network in frequency domain. In: *Proceedings of International Conference on Neural Information Processing*, 2019. 264–274
- 57 Oymak S, Soltanolkotabi M. Toward moderate overparameterization: global convergence guarantees for training shallow neural networks. *IEEE J Sel Areas Inf Theor*, 2020, 1: 84–105
- 58 Chan K H R, Yu Y D, You C, et al. ReduNet: a white-box deep network from the principle of maximizing rate reduction. 2021. ArXiv:2105.10446
- 59 E W. The dawning of a new era in applied mathematics. *Not Amer Math Soc*, 2021, 68: 565–571
- 60 Li Z L, You C, Bhojanapalli S, et al. Large models are parsimonious learners: activation sparsity in trained transformers. 2022. ArXiv:2210.06313
- 61 Zhang Z, Lin Y, Liu Z, et al. MoEfication: transformer feed-forward layers are mixtures of experts. In: *Proceedings of Findings of the Association for Computational Linguistics*, 2022. 877–890

Appendix A

Feature pathways in Figure 5. All eight feature pathways of the neurashed model are presented in Figure A1. Figure A1(a) and (b) correspond to Class 1 and Class 2, respectively.

Experimental details. In the experimental setup of the right panel of Figure A1(b), all amplification factors at initialization are set to independent uniform random variables on $(0, 0.01)$. We use $g^-(\lambda_F) = 1.022^{-\frac{1}{4}}\lambda_F$ and $g^+(\lambda_F) = 1.022^{\frac{1}{4}}\lambda_F$ for all hidden nodes except for the 7th (from left to right) node, which uses $g^+(\lambda_F) = 1.022^{\frac{3}{4}}\lambda_F$. In the early phase of training, the firing pattern on the second level improves at distinguishing the two types of samples in Class 1, depending on whether the 1st or 3rd node fires. This also applies to Class 2. Hence, the mutual information between the second level and the input tends to $\log_2 4 = 2$. By contrast, in the late stages, the amplification factors of the 1st and 3rd nodes become negligible compared with that of the 2nd node, leading to indistinguishability between the two types in Class 1. As a consequence, the mutual information tends to $\log_2 2 = 1$. The discussion on the first level is similar and thus is omitted.

Analyses for local elasticity. A common question from the reviews is whether neurashed can offer more quantitative results and insights into other phenomena. I have obtained the following new results.

Theorem A1. For a neurashed with m nodes in each level, select b_1 nodes uniformly at random at each level except for the last level; next, for each class, select b_2 nodes uniformly at random at each level except for the last level; last, for each class, draw edges from the $b_1 + b_2$ selected nodes to those in the consecutive levels. Let $g^+(x) = c + x$ and $g^-(x) = x$ for some $c > 1$. Suppose that each class has the same number of training examples. In the regime where the number of iterations tends to ∞ and m is sufficiently large, a backpropagation at Class j gives $\frac{\Delta Z_j}{\Delta Z_{j'}} \rightarrow \frac{b_1 + b_2}{b_1} > 1$.

Above, ΔZ_j denotes the change of the logit associated with Class j . The above inequality implies that an image has a more significant impact on images from the same class compared with those from other classes. This is a quantitative characterization of local elasticity. Next, regarding the effect of the number of layers, we have the following theorem.

Theorem A2. Under the same assumptions as the above theorem, if the number L of layers grows sufficiently quickly, then we have $\frac{\Delta Z_j}{\Delta Z_{j'}} \rightarrow \infty$.

Here we provide some intuition for the proof. The starting point is to use the asymptotic expansion $\Delta Z_j = (1 + o_{\mathbb{P}}(1))[(c(t + 1)b_2 + c(Kt + 1)b_1)^L - (ctb_2 + cKtb_1)^L]$. Besides, I now can show that neurashed can also shed light on self-supervised/contrastive

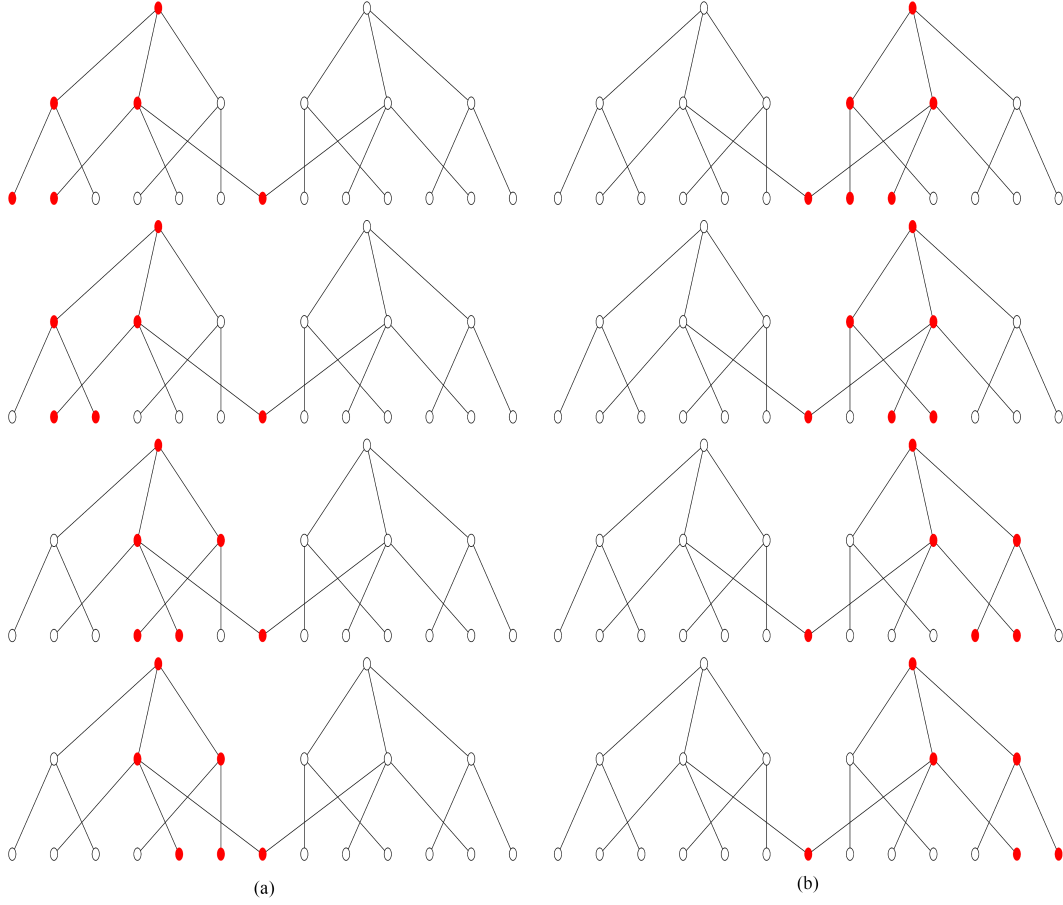


Figure A1 All feature pathways of the neurashed model in Figure 5. (a) Class 1; (b) Class 2.

learning. Roughly speaking, related inputs have similar feature pathways and so tend to be grouped together in new representations. More details follow in the final revision.

After N iterations, suppose there are N_1, N_2 , and N_3 times from the three classes, respectively. Then

$$f_j(x) = \begin{cases} a^2 N_1 (7N_1 + 5N_2 + 2N_3), & y = C_j, j = 1, \\ a^2 N_2 (5N_1 + 7N_2 + 2N_3), & y = C_j, j = 2, \\ 2a^2 N_3 (N_1 + N_2 + N_3) + aN_3 f', & y = C_j, j = 2, \\ 0, & y \neq C_j, \end{cases}$$

where f' denotes the contribution from F .

Now, suppose we update the neural networks using an example from C_1 . That is, replace N_1 by $N_1^+ := N_1 + 1$. Let us see how the prediction on C_2 and C_3 changes. Note that

$$\begin{aligned} & a^2 N_1^+ (7N_1^+ + 5N_2 + 2N_3) - a^2 N_1 (7N_1 + 5N_2 + 2N_3) \\ &= a^2 (N_1 + 1)(7N_1 + 5N_2 + 2N_3 + 7) - a^2 N_1 (7N_1 + 5N_2 + 2N_3) \\ &= 7a^2 N_1 + a^2 (7N_1 + 5N_2 + 2N_3 + 7). \end{aligned}$$

For C_2 we see

$$a^2 N_2 (5N_1^+ + 7N_2 + 2N_3) - a^2 N_2 (5N_1 + 7N_2 + 2N_3) = 5a^2 N_2,$$

and for C_3 we have

$$2a^2 N_3 (N_1^+ + N_2 + N_3) + aN_3 f' - 2a^2 N_3 (N_1 + N_2 + N_3) - aN_3 f' = 2a^2 N_3.$$

Note that in general $5a^2 N_2 > 2a^2 N_3$ unless N_3 is really large. However, to overcome this difficulty, we can say that for nodes in F^L , their amplifying factors are constant, meaning that they are all always on or never updated. For example, we can let the \mathcal{F}^L nodes be on if any of its children is on.