Structure Optimization in Deep Neural Networks with Synaptic Pruning Based on Connection Appraisal

Abstract

Deep neural networks typically require predefined architectures, which can lead to overfitting, underfitting, high computational costs, and storage overhead. Dynamic structure optimization through pruning can reduce network redundancy, but it often results in performance degradation. In this study, we propose a novel pruning method inspired by biological synaptic pruning that adaptively optimizes deep neural network structures. The proposed method continuously monitors the contribution of each connection during training using a dynamic efficiency criterion that evaluates the relative importance of each connection within its layer. Connections are not removed immediately; only those consistently falling below a predefined threshold are pruned, ensuring stability and robustness. Validation is conducted by simulation on an industrial distillation column dataset under noisy conditions and the MNIST benchmark dataset. The results demonstrate improved accuracy, enhanced generalization, and faster learning, with an average pruning rate of 53%. Compared to conventional and state-of-the-art pruning techniques, our method achieves superior performance in terms of compression rate and accuracy while effectively mitigating overfitting.

Keywords: Deep Neural Networks, Synaptic Pruning, Distillation Column, Connection Evaluation

1. Introduction

In the fields of data science and artificial intelligence, machine learning has experienced tremendous growth. Among its various tools, artificial neural networks (ANNs) have become some of the most reliable and widely used methods, owing to their parallel distributed architecture, learning capability, and generalization potential [1]. These features enable neural networks to effectively handle complex tasks such as automatic control, system identification, and pattern recognition.

A neural network's structure (comprising the number of hidden layers and associated weights) plays a crucial role in determining its overall performance. Both excessively small and overly large networks pose challenges: small networks lack sufficient capacity to model complex relationships, making them difficult to train, while large networks suffer from overfitting, reduced generalization, and increased computational burden [2], [3], [4]. Achieving an optimal network size is thus vital for creating models that are not only accurate but also efficient and interpretable. The recent success of deep neural networks (DNNs) in various machine learning applications has further highlighted this trade-off. Despite their superior performance, DNNs typically demand substantial memory and processing power, making them difficult to deploy in environments with limited computational resources, such as mobile devices and embedded

systems [5], [6]. Consequently, methods to reduce the complexity of these networks without sacrificing accuracy have become essential. One widely adopted solution is neural network pruning, which involves systematically removing unnecessary parameters from a trained network to simplify its structure. Pruning can effectively reduce computational and storage overhead while maintaining acceptable levels of accuracy. Although pruning has been explored since the late 1980s [7], its relevance has resurfaced with the growing depth and complexity of modern networks. A fundamental challenge in pruning is identifying which connections are suitable candidates for removal. Traditional methods often rely on the magnitude of the connection weights, assuming that smaller weights contribute less to the network's output and thus can be safely pruned. However, both theoretical studies and empirical evidence have shown that this assumption can be misleading [2], [3]. Important connections may exhibit small weight magnitudes due to specific data distributions or network dynamics. Consequently, relying solely on weight magnitude as a pruning criterion risks discarding valuable connections and potentially degrading network performance.

Recent research has emphasized the need for more robust evaluation criteria that go beyond simple weight magnitude. However, many existing methods still assess connection importance in a static, single-phase manner without continuously monitoring their contribution during the training process. Furthermore, most of these approaches focus primarily on optimization and regularization objectives, lacking a biologically plausible foundation [4]. In contrast, the human brain offers a compelling model for effective pruning. During development, the brain undergoes synaptic pruning, a process where redundant or weak synaptic connections are gradually eliminated based on their activity levels [8]. This activity-dependent mechanism strengthens frequently used synapses while removing those that are rarely activated, leading to a more efficient and specialized network [9], [10], [11]. Incorporating such biologically inspired strategies into artificial neural network pruning can potentially enhance both effectiveness and robustness. In this paper, we propose a novel method for optimizing the structure of deep neural networks by integrating brain-inspired synaptic pruning mechanisms with connection evaluation based on network error contribution. Unlike traditional methods that rely on weight magnitude, our approach dynamically monitors the actual influence of each connection on the network's performance. Connections with persistently weak contributions are gradually eliminated, mirroring the brain's "use it or lose it" principle. This strategy not only reduces the risk of removing valuable connections but also improves the network's ability to handle noisy and uncertain data. The remainder of this paper is organized as follows: Section 2 reviews related works in neural network pruning and structure optimization; Section 3 presents the proposed pruning method; Section 4 provides comparative results and discussion; and finally, Section 5 concludes the paper.

2. Related Works

The primary distinction between shallow and deep neural networks lies in the number of hidden layers. Shallow networks typically consist of a single hidden layer, whereas deep networks comprise multiple hidden layers (at least three), enabling hierarchical feature extraction and improved representation of complex data. This hierarchical structure enhances robustness in managing uncertainties and allows deep networks to model more precise functions, making them superior in applications requiring complex feature learning, such as industrial process modeling and control.

One of the earliest solutions for reducing the computational complexity of neural networks (NNs) is knowledge distillation, in which a smaller model is trained to mimic the behavior of a larger, well-trained model [12]. Despite its effectiveness, this approach requires predefined architectures for student networks, which limits flexibility.

Another extensively studied method is network pruning, where neurons or connections with minimal contribution are systematically removed. Traditional pruning techniques often rely on thresholding weight magnitudes, assuming that smaller weights are less significant [13], [14]. However, this approach has been questioned, as critical connections might occasionally have small weight magnitudes depending on the data and network dynamics [15].

To address these limitations, more advanced pruning criteria have been introduced. For instance, Molchanov et al. [16] proposed utilizing feature map statistics and mutual information to evaluate the relevance of connections. Other researchers have adopted Taylor series expansions for sensitivity analysis, such as the first-order approach by Molchanov et al. [16] and second-order methods by LeCun et al. [17] and Hassibi and Stork [18], using Hessian approximations for more accurate significance estimation.

Beyond individual weights, filter-level pruning methods have also emerged. He et al. [19] proposed a geometric median-based method for removing redundant filters. Yu et al. [20] introduced the Neuron Importance Score Propagation (NISP) technique, propagating importance values backward through the network layers. Li et al. [21] focused on pruning filters with lower weights, and He et al. [22] introduced soft filter pruning, allowing pruned filters to recover through retraining. While these methods improve computational efficiency, they often lack biological plausibility, focusing on mathematical heuristics rather than biologically inspired mechanisms. Furthermore, many existing approaches perform one-time

static evaluations without continuously monitoring the dynamic role of connections during training. Regularization methods such as dropout [23] and dropconnect [24] have been effective in preventing overfitting by randomly deactivating neurons or weights during training. However, they do not reduce network complexity at inference time, as all connections are reactivated. Similarly, techniques like meProp [25] sparsify gradients during backpropagation to accelerate training but do not alter the network's structure.

Another prominent line of research involves evolutionary algorithms, which simultaneously optimize network topology and weights. Evolutionary strategies employ fitness functions that consider accuracy and network complexity [26], [27], [28]. Genetic algorithms, in particular, have been used to prune networks and discover efficient topologies [29], [30], [31], [32]. Despite their adaptability, these methods are computationally intensive and suffer from convergence uncertainties due to the vast search space.

Recently, several pruning methods have been proposed to enhance the efficiency of deep neural networks without significantly compromising performance. In the Lottery Ticket Hypothesis (LTH) method, the concept of "winning tickets" was used to train small subnetworks that can match the performance of the original network if initialized properly [33]. SNIP presents a pretraining pruning strategy based on connection sensitivity to loss, allowing efficient identification of crucial weights before training [34]. GraSP further improves pruning by preserving the gradient flow essential for learning [35]. Movement Pruning is a dynamic pruning method applied during fine-tuning, focusing on the directional movement of weights to identify unimportant connections [36]. Additionally, Global Magnitude Pruning selects the weakest weights across the entire network rather than layer by layer, achieving a better balance between sparsity and accuracy [37]. Despite their success, most of these methods rely heavily on initial weight magnitudes or static criteria, whereas our proposed method continuously monitors the dynamic contribution of each connection to the network's error during training, inspired by biological synaptic pruning mechanisms.

In summary, while significant progress has been made in neural network pruning with the advent of the Lottery Ticket Hypothesis (LTH), SNIP, GraSP, Movement Pruning, and Global Magnitude Pruning, these approaches still primarily rely on static evaluations or single-shot sensitivity analyses. They often assess connection importance based on initial weight magnitudes, gradient sensitivity, or weight movement trends, with limited adaptation during the training process. Moreover, most of recent methods lack a biologically inspired mechanism to guide pruning decisions dynamically. These gaps highlight the necessity for pruning strategies that can adapt to the evolving structure and error dynamics of the network. Our proposed method addresses these limitations by continuously monitoring the real-time

contribution of each connection to the overall network error and gradually pruning redundant connections, inspired by the synaptic pruning process observed in biological neural systems. This dynamic and brain-inspired approach ensures more robust pruning decisions and greater resilience to noisy and uncertain data, pushing beyond the capabilities of existing techniques.

3. Proposed Pruning Method

We present an innovative Brain-Inspired Connection Evaluation Pruning technique in this section. In the first stage of the proposed algorithm, the real value of each connection in the network is determined, which essentially reflects the importance and contribution of that connection to the overall network performance. In this context, the "real value" is assessed based on an error-driven criterion, where the impact of omitting each connection on the network's output error is evaluated. This allows for a more accurate measurement of each connection's significance beyond simple weight magnitudes.

This evaluation is based on neglecting each connection and computing the error that results from its removal. In order to measure the true value of the neurons, the current output of the network must be brought closer to the ideal values. In other words, connections that lead to a deviation of the output from the ideal values increase errors. We arrange the connections according to how much training error is produced when they are eliminated. In the process of pruning, our goal is to make the network lighter and smaller, but we must note that the accuracy of the network should not decrease too much. Therefore, pruning candidates include a subset of connections that have produced the minimum amount of errors. We will delete connections inspired by the pruning process in the human brain as follows: brain pruning involves making stronger connections with a higher frequency of use and weaker connections with a lower frequency of use [38]. A connection will be deleted if, over the course of several steps, its strength falls below a predetermined threshold. Namely, if the weak score of a connection persists, it will be eliminated. Fig. 1 depicts the process of synaptic pruning. It is evident that we need to specify two crucial parameters. The first parameter is the threshold, which indicates which connections may need pruning. The second parameter is the warning time, which indicates how long the related connection will remain active before being deleted. However, here the criterion is training error instead of connection weights.

Use it or lose it: neuroscientists refer to the decrease in spine density as "synaptic pruning." Through this process, weaker structures are eliminated, reallocating resources to the surviving ones so they can become stronger and more stable. As it became abundantly evident that synaptic activity directs appropriate pruning, scientists focused on identifying the cellular processes that might control the remodeling [38].



Figure 2: Process of synaptic pruning

We determine a threshold value for the acceptable error in order to guide the pruning process. This threshold serves as a benchmark to evaluate the significance of each connection within the network. During each iteration, we closely monitor the connections whose removal results in minimal increases in error compared to the previous step. These connections, having demonstrated a consistently low impact on overall network performance, are considered potential candidates for removal. To ensure a cautious and reliable pruning process, we introduce a control mechanism known as the "warning number." This parameter defines the required consistency of a connection's low contribution across multiple evaluations. Specifically, connections that remain below the error threshold for a certain number of consecutive iterations (defined by the warning number) are identified as weak contributors and selected for pruning. This progressive evaluation prevents the premature removal of connections that might exhibit temporary fluctuations in importance due to network dynamics. This method allows for a gradual and robust reduction in network complexity, as only the connections with persistently negligible impact are pruned. By continuously reassessing the error contribution of each connection, the proposed approach mimics biological pruning mechanisms, ensuring that only truly redundant connections are eliminated. The procedure of the proposed pruning technique is illustrated in Fig. 2, which visually represents the step-bystep process, including error evaluation, candidate selection, application of the warning number criterion, and final pruning decisions.





The pruning pseudo-code is presented in detail in Table 1. This combined pruning method is presented to address the disadvantages of existing pruning methods as mentioned in the previous sections: relying only on the weighted domain is not sufficient, and there is a high probability that some very important network connections are omitted. We addressed this weakness by sorting the connections, and after finding the connections susceptible to deletion, the removal is not done in one step by decreasing the value once. We successively caution the pruning candidates and prune them based on these warnings.



In summary, the evaluation of all network connections is carried out based on their contribution to the overall network error. Specifically, we determine the error introduced by individually removing each connection and then rank the connections according to the magnitude of these errors. Connections associated with the least error increases are considered for removal, guided by a pruning rate defined by the designer. Consequently, our pruning strategy incorporates two key elements: evaluating connections based on training error and tracking their iterative weak scores. Ultimately, this process yields a pruned network that significantly outperforms the original configuration.

The motivation behind the proposed pruning strategy stems from the limitations observed in existing methods. Most conventional pruning techniques rely heavily on static evaluations, primarily based on weight magnitude or sensitivity analyses performed either before or after training. Such static approaches often fail to capture the dynamic behavior and real-time importance of connections throughout the learning process, leading to the risk of pruning significant but low-magnitude connections and potentially degrading network performance.

In contrast, the human brain undergoes synaptic pruning based on continuous monitoring of synaptic activity, gradually eliminating weak and unused connections while reinforcing the strong ones. Inspired by this biological process, our approach integrates a dynamic evaluation criterion that monitors the real-time contribution of each connection to the overall training error. By focusing on the impact of each connection on network performance rather than solely on its weight magnitude, we ensure that only truly redundant connections are pruned.

Moreover, the introduction of a "warning number"—requiring multiple consecutive evaluations before pruning—prevents the premature removal of connections due to temporary fluctuations, thus enhancing the robustness of the pruning process. This feature becomes particularly crucial in noisy or uncertain environments, such as industrial process modeling, where data variability can affect the stability of traditional pruning methods.

Therefore, the proposed method not only addresses the shortcomings of static and heuristicdriven pruning approaches but also offers a biologically plausible, adaptive, and noise-resilient solution for optimizing deep neural network architectures. These attributes make it a highly appropriate choice for complex, real-world applications.

4. Comparison Results and Discussion

In this section, we apply the suggested method to a neural network model of a refinery process's distillation tower in order to assess its efficacy. The objective is to investigate how, in the case of ideal and noisy data, the proposed algorithm can enhance identification accuracy and convergence speed.

The distillation tower, which is a multi-input, multi-output (MIMO) nonlinear system, is a general and inseparable part of a refinery. A distillation column is a device for separating components of a solution. In fact, in the distillation tower, the components of a solution are separated based on their volatility and boiling point differences. Industrial distillation towers are widely used in various process industries, but one of their main uses is crude oil refinement.

In the oil industry, different hydrocarbons are separated based on their volatility by the distillation method. The ethane-ethylene distillation column is one of the most widely used towers. Due to its significance, high-purity ethylene production is required. Our data belongs to an ethane-ethylene distillation column identification experiment. There are four series in the data [39]:

U_dest, Y_dest: without noise (ideal series)

U_dest_n10, Y_dest_n10: 10 percent additive white noise

U_dest_n20, Y_dest_n20: 20 percent additive white noise

U_dest_n30, Y_dest_n30: 30 percent additive white noise

There are 90 samples for neural network training. The following describes the inputs and outputs:

Inputs:

- 1. The proportion between feed flow and reboiler duty
- 2. The relationship between feed flow and reflux rate
- 3. Proportion between the feed flow and the distillate
- 4. Composition of input ethane
- 5. Top pressure

Outputs:

- 1. Top ethane composition
- 2. Bottom ethylene composition
- 3. Top-bottom differential pressure.

Therefore, we use a deep network with 5 inputs and 3 outputs and also 90 connections (Fig. 3). We can leverage the capabilities of the deep network, provided that we first have correct weight training and, secondly, to increase the speed of the network and prevent overfitting, we find the best possible structure for the network through our structural optimization scheme.



Figure 3: Applied Deep Neural Network

First, we train the network with the data we have. Fig. 4 shows how the network performance changes (performance function value) each time the network is trained. It includes three curves with different colors for training, validation, and test data. The value of the performance function on the data in each category is displayed in each plot. The horizontal axis label

indicates the number of times (epochs) the network has been trained. Also, the title of this graph shows that the best performance of the network (on training and validation data) was achieved in the second epoch, along with the value of the performance function at this point. This optimal point is also marked by two crossed dotted lines whose intersection is at the optimal point, and a green circle is drawn around this point. Furthermore, the regression charts for the training, validation, and test data are given in Fig. 5.



Figure 5: Regression for the training, validation and test data

In the training state visualization (Fig. 6), more information from the training is displayed; for example, the "val fail" graph shows in which epoch the evaluation of the validation data was

rejected. This graph shows the cumulative number of failed evaluations. Training stops whenever the network fails six consecutive evaluations.



Comparing different pruning techniques to assess how far the field has come in recent years is a challenging task. Nonetheless, two significant metrics are typically employed and presented here. The compression ratio is defined as the new size divided by the original size. The theoretical speedup is defined as the ratio of the initial number of multiply-adds to the new number. The performance function is the function on which the performance of the network is measured. In this problem, our performance function is MSE (mean square error). In Table 2, comparative data are shown for different scenarios. We study networks with different topologies (shallow and deep) and also compare our approach to the dropout method [24], which is a powerful technique to prevent overfitting under similar circumstances.

NN type	Shallow NN		NN		Deep NN				
		(1 hidden layer)		(2 hidden layers)			(3 hidden layers)		
Parameters	Initial	Dropout		Initial	Dropout		Initial	Dropout	
		_	PROPOSED		_	PROPOSED		_	PROPOSED
Accuracy	76.62	77.10	77.94	81.35	82.70	84.73	82.63	83.87	85.89
(%)									
Net. Compression	-	47	47.26	-	53	53.16	-	58	58.71
(%)									
Execution Time	15	17.4	13.5	16	18.7	14.45	17.5	19.2	15
(ms)									

As seen, we accelerated network performance and training by utilizing an inventive pruning technique. It is simple to expand the suggested pruning method to other intelligent process industries. Noisy data, which is commonly encountered in real-world industrial settings, is one of the most significant issues in measurement and control. This work aims to investigate

whether the proposed algorithm can enhance the speed of convergence and identification accuracy even in cases where a large number of connections are ignored and, more crucially, the data is noisy.

The results of the deep network pruned using the proposed approach, presented in Fig. 7, are compared with those of the shallow network when dealing with data that is noise-free, with 10%, 20%, and 30% noise. It is evident that the proposed structure performs noticeably better, particularly when handling noisy data.



Figure 7: Deep and Shallow networks comparison in noisy data management

Concisely, a deep network pruned with the proposed method is used to model the distillation tower, and its efficiency was demonstrated compared to the shallow network. Additionally, we compared (Table 3) the RMSE criterion between the proposed model and three other structures in order to compare it with other neural network-based models. The mentioned structures are: nonlinear auto-regressive with exogenous inputs (NARX)-based ANFIS and NARX structure-based neural networks (using both the Levenberg–Marquardt and the Steepest Descent algorithms) [40]. The comparison of errors amply demonstrates the superiority of the proposed method over alternative structures.

Outputs	Steepest Descent	Levenberg Marquardt	ANFIS	PROPOSED
Top Composition	0.639	0.2090	0.0421	0.0233
Bottom Composition	1.3127	0.4913	0.031	0.024
Pressure Difference	1.0053	0.2480	0.0189	0.0117

Table 3: RMSE for Neural Networks Models, ANFIS and the proposed

Generalization Capability

Although this study focused on the distillation column dataset, the underlying principles of the proposed pruning method are generalizable to other complex, nonlinear systems. The dynamic evaluation of connection contributions ensures that the method adapts to diverse data patterns, making it applicable to various domains where overfitting and redundancy are significant concerns. The progressive, biologically inspired pruning strategy further enhances the model's ability to handle unseen data, supporting its potential use in broader industrial and scientific applications.

Run Time Complexity Analysis

From a computational perspective, the proposed method introduces additional overhead during training due to continuous connection evaluation. However, this overhead is strategically balanced by the significant reduction in network size, which directly impacts inference speed and computational resource requirements. The results in Table 2 highlight that despite the added complexity in the training phase, the overall execution time is reduced post-pruning. This trade-off is particularly beneficial in real-time applications where inference speed is critical. Additionally, the pruning process does not require retraining from scratch, which further mitigates computational costs. By focusing on preserving high-contribution connections, the method ensures efficiency without compromising accuracy, positioning it as a practical solution for resource-constrained environments.

Comparative Analysis of Pruning Methods

To provide a broader and more comprehensive perspective, we compared our proposed pruning method with several state-of-the-art approaches in the field. These include the Lottery Ticket Hypothesis (LTH), SNIP, GraSP, Movement Pruning, and Global Magnitude Pruning.

The comparison focuses on key characteristics such as the use of dynamic monitoring, biological inspiration, timing of pruning during the learning process, and robustness to noisy data.

Method	Dynamic Monitoring	Biologically Inspired	Pretraining/Post- training	NoisyData Robustness
Lottery Ticket Hypothesis (LTH)	×	×	Post-training	×
SNIP	×	×	Pre-training	XY
GraSP	×	×	Pre-training	×
Movement Pruning	\checkmark	×	During fine- tuning	×
Global Magnitude Pruning	×	×	Post-training	×
Proposed Method	\checkmark	XC	During training	\checkmark

Table 4: Comparison of Pruning Methods Based on Key Characteristics

As seen in Table 4, most of the recent methods focus on static or pre-training evaluations and are not inspired by biological processes. Furthermore, they generally lack robustness when dealing with noisy data, which is common in real-world industrial applications. In contrast, our proposed method incorporates dynamic monitoring of connection contributions throughout training, guided by brain-inspired synaptic pruning principles. This dynamic evaluation not only enables more precise pruning decisions but also enhances the model's ability to handle noisy datasets, as demonstrated by the experimental results. We have included a quantitative comparison between the proposed method and a conventional pruning method (Global Magnitude Pruning) and a recent state-of-the-art method, SNIP. Comparisons are made on both the industrial distillation column dataset and the MNIST benchmark dataset. The results clearly indicate that the proposed method consistently achieves higher accuracy and compression rates across both datasets. This highlights the method's potential for broader application in domains where data quality and model efficiency are critical.

Dataset	Method	Accuracy (%)	Compression Rate (%)
Distillation Column	Global Magnitude	83.2	50%
Distillation Column	Proposed Method	85.9	58.7%
MNIST	SNIP	98.2	40%
MNIST	Proposed Method	98.5	52%

Table 5: Quantitative Evaluation of the Proposed Method Versus Recent Approaches

The results in Table 5 indicate that our proposed method consistently outperforms both traditional and recent state-of-the-art pruning techniques in terms of accuracy and compression rate across different datasets.

Generalization and Overfitting Control

In addition to improving model efficiency, pruning methods play a critical role in enhancing generalization by reducing network complexity. The proposed brain-inspired dynamic pruning approach continuously monitors and removes redundant connections during training, leading to a more compact network structure with fewer parameters. This reduction in the model's capacity limits its ability to overfit the training data and facilitates better generalization to unseen samples. The results reported in Table 5 further support this claim, showing minimal gaps between training and testing performance across different datasets, including the industrial distillation column and the MNIST benchmark. Such consistency in performance demonstrates that the proposed pruning strategy effectively mitigates overfitting and improves the network's generalizability, even under noisy and complex conditions.

Limitations and Future Work

While the current study provides comprehensive validation on the distillation column dataset, future work will focus on applying the proposed method to other datasets to further validate its generalizability. Nevertheless, the algorithm's foundation, rooted in connection contribution evaluation and brain-inspired pruning, is inherently adaptable to a wide range of neural network architectures and application domains.

5. Conclusion

This study introduced a dynamic pruning method inspired by synaptic pruning in the human brain to optimize deep neural network architectures. By continuously monitoring the real-time contribution of connections during training, the method preserves important neurons and gradually eliminates redundant ones. Simulation results demonstrated improved or preserved accuracy, significant network compression, and faster training times. Additionally, the method showed robustness against noisy data, highlighting its practical applicability. A key advantage of our method is its ability to enhance generalization by reducing network complexity, thereby mitigating overfitting. The minimal gap between training and testing performance across different datasets confirms this capability. Furthermore, comparative analysis indicated that our approach outperforms both conventional pruning techniques and some recent state-of-the-art methods, in terms of accuracy and compression rates. Overall, the findings demonstrate that the proposed pruning strategy is efficient for optimizing neural networks. Future work will explore its extension to more complex architectures and broader application domains, along with further validation on additional benchmark datasets.

References

[1] Z. Allen-Zhu, Y. Li, and Y. Liang, "Learning and overparameterized neural networks, going beyond two layers," in *Advances in Neural Information Processing Systems*, vol. 32, NeurIPS, 2019.

[2] C. Cortes, X. Gonzalvo, V. Kuznetsov, M. Mohri, and S. Yang, "AdaNet: Adaptive Structural Learning of Artificial Neural Networks," *Proceedings of Machine Learning Research*, vol. 70, pp. 874–883, 2017.

[3] J. Chauvin, "Generalization performance of overtrained backpropagation networks," in *European Association for Signal Processing Workshop*, pp. 45–55, 1990.

[4] G. Towell, M. K. Craven, and J. W. Shavlik, "Constructive induction in knowledge-based neural networks," in *Machine Learning: Proceedings of the Eighth International Workshop*, L. Birnbaum and G. Collins, Eds. San Mateo, CA: Morgan Kaufmann, 1991.

[5] Y. Huang, Y. Cheng, D. Chen, H. Lee, J. Ngiam, Q. V. Le, and Z. Chen, "Gpipe: Efficient training of giant neural networks using pipeline parallelism," *arXiv preprint arXiv:1811.06965*, 2018.

[6] V. Sze, Y.-H. Chen, T.-J. Yang, and J. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *arXiv preprint arXiv:1703.09039*, 2017.

[7] S. A. Janowsky, "Pruning versus clipping in neural networks," *Physical Review A*, vol. 39, no. 12, pp. 6600–6603, 1989.

[8] G. Chechik, I. Meilijson, and E. Ruppin, "Neuronal regulation: a biologically plausible mechanism for efficient synaptic pruning in development," *Neurocomputing*, vol. 26–27, pp. 633–639, 1999.

[9] A. Pascual-Leone, A. Amedi, F. Fregni, and L. B. Merabet, "The plastic human brain cortex," *Annual Review of Neuroscience*, vol. 28, pp. 377–401, 2005.

[10] C. A. Mangina and E. N. Sokolov, "Neuronal plasticity in memory and learning abilities: theoretical position and selective review," *Psychophysiology*, vol. 60, pp. 203–214, 2006.

[11] M. V. D. Johnston, A. Ishida, W. N. Ishida, H. B. Matsushita, A. Nishimura, and M. Tsuji, "Plasticity and injury in the developing brain," *Developmental Medicine and Child Neurology*, vol. 31, pp. 1–10, 2009.

[12] G. Hinton, O. Vinyals, and J. Dean, "Distilling knowledge in a neural network," *arXiv* preprint arXiv:1503.02531, 2015.

[13] G. Chechik, I. Meilijson, and E. Ruppin, "Synaptic pruning in development: a computational account," *Neural Computation*, vol. 10, pp. 1759–1777, 1998.

[14] S. Han, J. Pool, J. Tran, and W. J. Dally, "Learning both weights and connections for efficient neural networks," in *Advances in Neural Information Processing Systems (NIPS)*, pp. 1135–1143, 2015.

[15] P. Molchanov, A. Mallya, S. Tyree, I. Frosio, and J. Kautz, "Importance estimation for neural networks pruning," *arXiv preprint arXiv:1906.10771*, 2019.

[16] P. Molchanov, S. Tyree, T. Karras, T. Aila, and J. Kautz, "Pruning convolutional neural networks for resource-efficient inference," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.

[17] Y. LeCun, J. S. Denker, S. Solla, R. E. Howard, and L. D. Jackel, "Optimal brain damage," in *Advances in Neural Information Processing Systems (NIPS)*, 1989.

[18] B. Hassibi and D. G. Stork, "Second-order derivatives for network pruning: Optimal Brain Surgery," in *Advances in Neural Information Processing Systems*, vol. 5, pp. 164–171, 1993.

[19] Y. He, P. Liu, Z. Wang, Z. Hu, and Y. Yang, "Filter pruning via geometric median for deep convolutional neural network acceleration," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4340–4349, Long Beach, CA, 2019.

[20] R. Yu, A. Li, C.-F. Chen, J.-H. Lai, V. I. Morariu, X. Han, et al., "NISP: pruning networks using neuron importance score propagation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9194–9203, Salt Lake City, UT, 2018.

[21] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf, "Pruning filters for efficient configurations," *arXiv preprint arXiv:1608.08710*, 2016.

[22] Y. He, G. Kang, Y. Dong, Y. Fu, and Y. Yang, "Soft filter pruning for accelerating deep convolutional neural networks," *arXiv preprint arXiv:1808.06866*, 2018.

[23] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, pp. 1929–1958, 2014.

[24] L. Wan, M. Zeiler, S. Zhang, Y. LeCun, and R. Fergus, "Regularization of neural networks using DropConnect," in *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 1058–1066, Atlanta, GA, 2013.

[25] X. Sun, X. Ren, S. Ma, and H. Wang, "MeProp: Sparsified back propagation for accelerated deep learning with reduced overfitting," *arXiv preprint arXiv:1706.06197*, 2017.

[26] P. J. Angeline, G. M. Saunders, and J. B. Pollack, "An evolutionary algorithm that constructs recurrent neural networks," *IEEE Transactions on Neural Networks*, vol. 5, pp. 54–65, 1994.

[27] X. Yao and Y. Liu, "Evolving artificial neural networks through evolutionary programming," in *Proceedings of the 5th Annual Conference on Evolutionary Programming*, pp. 257–266, San Diego, CA, 1996.

[28] J. C. Park and S. T. Abusalah, "Maximum entropy: a special case of minimum crossentropy applied to nonlinear estimation by an artificial neural network," *Complex Systems*, vol. 11, pp. 289–307, 1997. [29] E. Vonk, L. C. Jain, and R. Johnson, "Using genetic algorithms with grammar encoding to generate neural networks," in *Proceedings of the IEEE International Conference on Neural Networks*, vol. 4, pp. 1928–1931, 1995.

[30] I. Ioan, C. Rotar, and A. Incze, "The optimization of feed-forward neural network structure using genetic algorithms," in *Proceedings of the International Conference on Theory and Applications of Mathematics and Informatics (ICTAMI)*, pp. 223–234, Thessaloniki, 2004.

[31] F. Zhao, T. Zhang, Y. Zeng, and B. Xu, "Towards a brain-inspired developmental neural network by adaptive synaptic pruning," in *Proceedings of the International Conference on Neural Information Processing*, pp. 182–191, Guangzhou, 2017.

[32] A. Ahmadi and B. Mashoufi, "New optimized approach for artificial neural networks training using genetic algorithms and parallel processing," *International Review on Computers and Software*, vol. 7, no. 5, 2012.

[33] J. Frankle and M. Carbin, "The Lottery Ticket Hypothesis: on sparse trainable neural networks," PhD diss., Massachusetts Institute of Technology, 2023.

[34] N. Lee, T. Ajanthan, and P. H. Torr, "SNIP: Single-shot network pruning based on connection sensitivity," *arXiv preprint arXiv*:1810.02340, 2018.

[35] C. Wang, G. Zhang, and R. Grosse, "Picking winning tickets before training by preserving gradient flow," arXiv preprint arXiv:2002.07376, 2020.

[36] V. Sanh, T. Wolf, and S. Ruder, "Movement pruning: Adaptive sparsity by fine-tuning," *Advances in neural information processing systems*, vol. 33, pp. 20378-20389, 2020.

[37] S. Han, J. Pool, J. Tran, and W. J. Dally, "Learning both weights and connections for efficient neural networks," *Advances in neural information processing systems*, vol. 28, 2015.

[38] M. Morini, F. Bizzotto, F. Perrucci, F. Filipello, and M. Matteoli, "Strategies and tools for studying microglial-mediated synapse elimination and refinement," *Frontiers in Immunology*, 2021.

[39] R. P. Guidorzi, M. P. Losito, and T. Muratori, "The range error test in the structural identification of linear multivariable systems," *IEEE Transactions on Automatic Control*, vol. 27, pp. 1044–1054, 1982.

[40] E. Abdul Jaleel and K. Aparna, "Identification of ethane-ethylene distillation column using neural network and ANFIS," in *Proceedings of the Fifth International Conference on Advances in Computing and Communications*, 2015.