

Are Sparse Neural Networks Better Hard Sample Learners?

Qiao Xiao¹

q.xiao@tue.nl

Boqian Wu^{2,3}

b.wu@utwente.nl

Lu Yin⁴

l.yin@surrey.ac.uk

Christopher Neil Gadzinski³

christopher.gadzinski@uni.lu

Tianjin Huang⁵

t.huang2@exeter.ac.uk

Mykola Pechenizkiy¹

m.pechenizkiy@tue.nl

Decebal Constantin Mocanu^{1,3}

decebal.mocanu@uni.lu

¹ Eindhoven University of Technology
The Netherlands

² University of Twente
The Netherlands

³ University of Luxembourg
Luxembourg

⁴ University of Surrey
United Kingdom

⁵ University of Exeter
United Kingdom

Abstract

While deep learning has demonstrated impressive progress, it remains a daunting challenge to learn from hard samples as these samples are usually noisy and intricate. These hard samples play a crucial role in the optimal performance of deep neural networks. Most research on Sparse Neural Networks (SNNs) has focused on standard training data, leaving gaps in understanding their effectiveness on complex and challenging data. This paper’s extensive investigation across scenarios reveals that most SNNs trained on challenging samples can often match or surpass dense models in accuracy at certain sparsity levels, especially with limited data. We observe that layer-wise density ratios tend to play an important role in SNN performance, particularly for methods that train from scratch without pre-trained initialization. These insights enhance our understanding of SNNs’ behavior and potential for efficient learning approaches in data-centric AI. Our code is publicly available at: https://github.com/QiaoXiao7282/hard_sample_learners.

1 Introduction

In the last decade, deep learning has seen remarkable developments, primarily benefiting from the increasing training data and the accompanying larger models [39, 43]. This progression, however, comes with high computational costs and complex optimization challenges. Recent insights reveal that not all training samples are equally important, with a small subset

contributing most to the loss [26, 56]. Additionally, studies [6, 62, 85, 63] have shown that eliminating redundant data and prioritizing training samples based on inforatic complexity can improve efficiency without sacrificing performance.

Investigating the data difficulty not only helps to improve the training efficiency, but also helps us to understand the principles that govern how deep models process data [2, 6, 17]. Recent studies have demonstrated that training with more challenging samples can lead to improved generalization [20, 54], offering valuable insights into the intricacies of deep learning model behavior. Building on the concept of such a training paradigm, researchers have also started experimenting with the addition of perturbations to input data to make it more challenging. In a related vein, the focus has shifted towards constructing adversarial samples, which intentionally perturb inputs to confuse deep learning models, aiming to enhance model safety and robustness [9, 9]. Concurrently, the introduction of data corruption is being employed to further advance the comprehension of how deep learning models can be trained more effectively and securely [44, 50].

However, learning from challenging samples, which are often complex or difficult for model training, can introduce spurious features [23, 48], increasing the risk of overfitting, particularly when training data is limited. Sparse Neural Networks (SNNs) [3, 12, 29, 57], known for efficiently eliminating redundant weights, have shown potential in mitigating overfitting [21, 25]. While the study of SNNs has primarily focused on standard training datasets, where they have demonstrated their ability to maintain performance with reduced computational cost, their behavior under more challenging conditions remains less thoroughly explored. This leads to a natural question: *Can SNNs perform well when trained on challenging samples?* Studies by [41] and [9] indicate that SNNs can reduce overfitting when trained with adversarial samples in standard data volumes. Varma T et al. suggest the Lottery Ticket Hypothesis (LTH) [14] is effective in smaller data volumes with extensive augmentation [52]. Additionally, He et al. [17] observe the sparse double descent phenomenon in network pruning, which is caused by model sparsity when addressing overfitting issues. However, these studies often focus on specific sparse models or isolated scenarios. This paper aims to answer a broader question: *Given the diverse methods for achieving sparsity, how do SNNs perform with challenging samples across different scenarios, and what factors contribute to their efficacy in these contexts?*

To answer this question, we undertake the exploration through the following avenues: (1) Samples with intrinsic complexity, where we identify challenging samples using EL2N (Error L2 Norm) [42] to assess learning difficulty, and (2) Samples with external perturbations, involving adversarial examples that subtly alter inputs to significantly affect model performance, and more noticeable corruptions like Gaussian noise and blurring to increase the sample difficulty. Through comprehensive experiments covering a wide range of sparse methods, model sizes and datasets, our study unveils several nuanced and occasionally surprising findings.

- We systematically analyze the effectiveness of model sparsification on difficulty sample training, considering the various conceptions of sample difficulty defined above. Our findings indicate that most SNNs can achieve or even surpass the accuracy of dense models at certain levels of sparsity.
- We extend our investigation to scenarios with limited training data and find that, in most cases, SNNs can achieve performance improvements over their dense counterparts, even at high levels of sparsity, when trained on challenging samples characterized by intrinsic complexity and external perturbations.

- Our findings suggest that the layer-wise density ratio in SNNs may contribute to the performance improvement in challenging training scenarios. In particular, maintaining a higher density in shallower layers positively impacts performance, especially in methods that train from scratch without pre-trained initialization.

2 Related work

2.1 Sparse Neural Networks

Dense to sparse. Dense to sparse training methodologies start with a dense model, and then strategically eliminate weights in one or several stages, interspersed with retraining for accuracy recovery. A typical method is Gradual Magnitude Pruning (GMP), which iteratively sparsifies weights based on their absolute magnitude, either globally or per layer, across several training steps [16, 18, 67]. Augmenting this, second-order pruning methods incorporate second-order information to potentially enhance the accuracy of the pruned models [61, 58, 65]. Unlike previous methods, retraining starts from a fully or well-trained dense model, the Lottery Ticket Hypothesis (LTH), restarts from the initial model state [12] or by rewinding to an earlier stage of the model [7, 15, 45] based on the given binary mask. In contrast to the aforementioned methods, sparsification can also be achieved from dense models early in training, before the main training phase, based on certain salience criteria [29, 63, 69].

Sparse to sparse. Unlike dense-to-sparse methodologies, where sparsification heavily relies on dense models, sparse-to-sparse training usually starts with a randomly initialized sparse topology before training. This training paradigm, starts with the initialization of a sparse subnetwork, followed by either maintaining its connectivity statically [36] or periodically searching for the optimal sparse connectivity through prune-and-regrow strategy [31, 67, 62] during training. For prune-and-regrow strategy, there exist numerous pruning criteria in the literature, including magnitude-based pruning [13, 61, 60], weight-balanced pruning [57], and gradient-based pruning [40, 64]. On the other hand, the criteria used to regrow weights back include randomness [37, 58], momentum [12], and gradient [13, 24, 31].

2.2 Learning on Hard Samples

Numerous studies have sought to define sample difficulty, shedding light on how deep neural networks evolve their data processing capabilities during training. It has been shown that deep learning models tend to learn difficult data later in the training process [11, 9]. Furthermore, training with hard samples has been found to accelerate the optimization of deep learning networks under enough training data volumes [47, 49]. Meanwhile, in the pursuit of learning efficiency, it has been demonstrated that training with a harder subset can maintain final performance [6, 54, 35].

On the other hand, deep neural networks (DNNs) are susceptible to malicious attacks, where specially perturbed inputs, known as adversarial samples, are crafted to challenge these models and train with them to improve their robustness [8, 46, 52, 53]. However, this training method has been observed to result in substantial robust generalization gaps, a phenomenon known as robust overfitting. Recently, sparse models have been proven to achieve better robust generalization [9, 41] and prevent overfitting problems [21] while achieving

more efficient training. However, these observations are under standard training data volumes, utilizing the full training dataset.

Unlike prior work, this paper systematically evaluates the effectiveness of SNNs learning on challenging samples, which are defined across a broader spectrum of scenarios. We also compare different sparsification methods and extend the exploration under reduced training data volumes.

3 Methodology and Evaluation

3.1 Sparse Neural Networks

In this paper, we primarily focus on unstructured sparsity, as these methods have been extensively studied in the literature, benefit from established benchmarks, and provide an optimal trade-off between accuracy and compression. To have a unified framework for SNNs, we use binary masks to simulate the implementation of model sparsity. Given a dense network with parameters $\theta_l \in \mathbb{R}^{d_l}$, where d_l is the dimension of the parameters in each layer $l \in \{1, \dots, L\}$, the sparse neural networks can be facilitated as $\theta_l \odot \mathbf{M}_l$, where $\mathbf{M}_l \in \{0, 1\}^{d_l}$ donates layer-wise binary mask, and \odot is the elementwise product. The sparsity ratio is determined by the fraction of weights set to zero, calculated as $s = 1 - \sum_l \|\mathbf{M}_l\|_0 / \sum_l d_l$. We choose the following representative sparsity methods for analysis:

- *Gradual Magnitude Pruning (GMP)*, as introduced in [16, 67], starting from a dense model, progressively sparsifies networks from dense model during training, by using the weight magnitude as the criterion for sparsity.
- *Lottery Ticket Hypothesis (LTH)* proposed by [14] is another commonly used sparsity method. It iteratively employs magnitude pruning during training to create binary masks and then re-trains using weights from step t . In our experiments, we set $t = 0$, which means we re-train with the initialized weights.
- *Magnitude After Training (OMP)*, which follows dense model training on a specific task, facilitates one-shot pruning using weight magnitude as the criterion, and then re-trains the model using the full learning rate schedule, we follow the setting as in [82].
- *SNIP* [49] is a typical prior-training pruning technique that globally removes weights with the lowest connection sensitivity score defined by $|\theta| \cdot |\nabla_{\theta} \mathcal{L}|$, and keeps the sparse topology of the model fixed throughout training.
- *Sparse Evolutionary Training (SET)*, a pioneering method for dynamic sparse training proposed by [57], begins with an initially sparse subnetwork and concurrently updates its topology and weights during training through a dynamic prune-and-grow strategy.

The main implementation setup for SNNs primarily follows [13, 51]. Further details can be found in Appendix A.1.

3.2 Experiments on Samples with Intrinsic Complexity

In this section, we evaluate the performance of dense and SNN models trained on samples with intrinsic complexity, as measured by EL2N scores. We first train a model on the entire dataset to calculate these scores and then classify the top 50% of samples with the highest scores as hard samples. Subsequently, we retrain the models from scratch on this challenging subset. We conduct experiments using ResNet18 for CIFAR100 and ResNet34 for TinyImageNet. Details on model training and sparsity are provided in Appendix A.2.




Figure 1: Examples of five CIFAR-100 training images for two randomly selected classes (apple and bus), showcasing those with the higher and lower EL2N scores. Images with lower scores typically feature simpler backgrounds and clear objects, whereas those with higher scores frequently display complex backgrounds or color biases.

3.2.1 EL2N Scores Based Measurements

The EL2N (Error L2 Norm) score, as proposed by [42], is defined as the L2 distance between the model predicted probability and the one-hot label of the sample. Given a training sample (x, y) , the EL2N score is defined as: $\mathbb{E} \|p(x; \theta) - y\|_2$, where $p(x; \theta) = \sigma(f(x; \theta))$ denotes the neural network output in the form of a probability vector, and σ is the softmax function.

Intuitively, samples with higher EL2N scores are often seen as more challenging and are therefore categorized as having intrinsic complexity in our paper. In Figure 1, we provide examples of samples with both higher and lower EL2N scores to illustrate how they help in identifying samples with intrinsic complexity.




Figure 2: Comparison of dense and SNNs models trained with EL2N score-filtered samples across CIFAR100 and TinyImageNet, with sparsity ratios from 10% to 90%. Sub-figures (a) and (b) display results trained with top 50% filtered samples, while sub-figures (c) and (d) show results from the top 30% filtered samples.

Results and Analysis: We observe that most SNNs methods are able to consistently match or even surpass dense models when trained on samples with intrinsic complexity on both CIFAR-100 and TinyImageNet datasets at a data ratio of 0.5. Specifically, as illustrated in Figure 2 (a) and (b), methods like SET or SNIP demonstrate superior performance compared to dense models, especially at higher sparsity levels. However, the LTH method exhibits a decline in performance, which suggests that when trained with hard samples, the masks derived from pre-trained models might not be optimal for the LTH approach.




Figure 3: The comparison covers sparsity ratios and data ratios ranging from 10% to 90% on CIFAR-100 dataset using ResNet18.

At low training data volumes (e.g. data ratio=0.3), most SNNs tend to offer more advantages over their dense counterparts, especially when sparsity levels are higher. From Figure 2 (c) and (d), we can find that SET and SNIP considerably outperform their dense counterparts, especially at higher sparsity levels, when trained with only 30% of the harder data. Moreover, to be more specific, consider the SET method as an example. As shown in Figure 3, when trained on a reduced dataset size, such as only 20% of the training data, SNNs trained with SET can significantly outperform dense models in test accuracy, particularly at higher sparsity levels. This suggests that training with less data may more easily lead to overfitting in dense models, potentially degrading their performance.

3.3 Experiments on Samples with External Perturbing

In this section, we will introduce two different types of data perturbations for training datasets. The first type consists of common corruptions such as Gaussian noise, blurring, or other visible image degradations. The second type involves adversarial attacks, which are imperceptible to the human eye but can substantially impact model performance.

3.3.1 Samples with Common Corruptions

We evaluate the performance of SNNs and dense models trained on samples impacted by common image corruptions, which introduce visible distortions that shift the data distribution from the original dataset. These distortions, often encountered in real-world applications, include Gaussian noise, impulse noise, and defocus blur. Following the methodology in [22], we applied these image corruptions to the datasets for a more comprehensive assessment. We conduct experiments using ResNet18 and VGG19 for CIFAR-100, and ResNet34 for TinyImageNet, with further details on model training and sparsity provided in Appendix A.3.




Figure 4: Comparison of dense and SNNs training on samples with common corruptions across CIFAR100 and TinyImageNet datasets with sparsity ratios ranging from 10% to 90%. The sub-figures (a) and (b) showcase experiments conducted on full data volume, while the last two (c) and (d) are conducted on a 30% data ratio.

Results and Analysis: In our experiments, Figure 4 presents the main results with these corruptions applied at severity level 5, with additional results for levels 2, 4 and 6 provided in Appendix B. We observe that SNN methods can perform comparably to or even surpass dense models when trained with samples affected by common corruptions on both the CIFAR-100 and TinyImageNet datasets at certain sparsity ratios. Specifically, under full training data conditions, as shown in Figure 4 (a) and (b), most SNNs methods can outperform dense models mainly at lower sparsity ratios. This is likely because SNNs at higher sparsity ratios may have a reduced capacity to learn from a large number of challenging examples. Meanwhile, the LTH method requires a lower sparsity ratio to maintain decent performance, as the masks obtained from challenge samples pre-training models may not be

ideally suited for the LTH methodology, the finding is consistent with the previous observation.

At lower training data volumes (e.g., data ratio = 0.3), SNNs generally demonstrate greater benefits compared to their dense counterparts, particularly at higher sparsity levels. In particular, SET and SNIP outperform their dense model, especially at higher sparsity levels, as shown in Figure 4 (c), (d), and Figure 5 (b). This performance enhancement is likely due to SNNs’ ability to mitigate overfitting issues when training with challenging samples under limited data conditions, a capacity that has been demonstrated in other studies [24, 25].




Figure 5: Comparison of dense models and SNNs trained with samples with common corruptions using VGG19 on CIFAR-100, under full data volume (a) and 30% data volume (b). This comparison spans a range of sparsity ratios, from 10% to 90%.

3.3.2 Samples with Adversarial Attack

PGD Attack and PGD Adversarial Training: Adversarial samples are special instances perturbed by well-designed changes with the purpose of confusing deep learning models. We conduct experiments using the well-known Projected Gradient Descent (PGD) attack method [63] for generating adversarial samples. The perturbations at $t + 1$ can be defined as follows:

$$\delta^{t+1} = \text{proj}_{\mathcal{P}} [\delta^t + \alpha \cdot \text{sgn}(\nabla_x \mathcal{L}(f(x + \delta^t; \theta), y))] \quad (1)$$

with a step size α , where \mathcal{P} is the set $\delta: \|\delta\|_p \leq \epsilon$, and the ℓ_p norm of the perturbation δ is constrained to a small constant ϵ . During training, the optimization problem is transformed into a min-max problem: $\min_{\theta} \mathbb{E}_{(x,y) \in \mathcal{D}} \max_{\|\delta\|_p \leq \epsilon} \mathcal{L}(f(x + \delta; \theta), y)$, where $f(x; \theta)$ is a network parameterized by θ , and the input data $(x, y) \in \mathcal{D}$ are combined with the perturbation δ to generate adversarial samples, which are then used to minimize the empirical loss function \mathcal{L} .

Our experiments involve two popular architectures, VGG-16 and ResNet-18, evaluated on CIFAR-10 and CIFAR-100, respectively. We assess both adversarial accuracy on perturbed test data and clean accuracy on unperturbed datasets. The performance is evaluated using the final checkpoint after training completion. Further details on model training and sparsity are provided in Appendix A.3.




Figure 6: Comparison of clean and adversarial test accuracy between dense models and various SNNs methods on CIFAR-10 with VGG16 and CIFAR-100 with ResNet18 at overall sparsity levels of 0.9. Sub-figures (a) and (b) are models trained on full data volume, (c) and (d) are models trained using only 50% of the training data.

Results and Analysis: When trained with adversarial attack samples at full data volume, SNNs consistently outperform dense models in clean and adversarial combined accuracy. As shown in Figure 6 (a) and (b), it is evident that at a sparsity level of 90%, most SNNs for VGG16 on CIFAR-10 and ResNet18 on CIFAR-100 maintain comparable clean accuracy and exhibit superior adversarial accuracy. This is consistent with findings from [9], which suggest that SNNs can mitigate overfitting issues when training with adversarial attack samples. Among the SNNs methods, SET slightly stands out in performance across both datasets slightly. Specifically, for ResNet18 on CIFAR-100, SET demonstrates superior results, while LTH and OMP show slightly lower adversarial accuracy compared to other SNN methods.

At low training data volumes, using only 50% of the training data, SNNs can also outperform the dense model in terms of clean and adversarial combined accuracy in most cases. As in Figure 6 (c) and (d), at a sparsity level of 90%, where the adversarial accuracy of SNNs surpasses that of dense models, a trend consistent with the full data regime. More results on other sparsity ratios can be found in Appendix C.

4 Empirical Analysis and Discussion

Recognizing the benefits of model sparsity in various challenging sample scenarios, this section will empirically explore the factors contributing to performance improvements in sparse neural networks when training with hard samples.

4.1 Which Layers Are Getting Sparsified?

In Sparse Neural Networks (SNNs), the layer-wise density ratio indicates the proportion of non-zero parameters in each layer. Given that SNNs perform differently when trained with challenging samples, our initial investigation aims to explore how the layer-wise density ratio is distributed across various SNN methods.

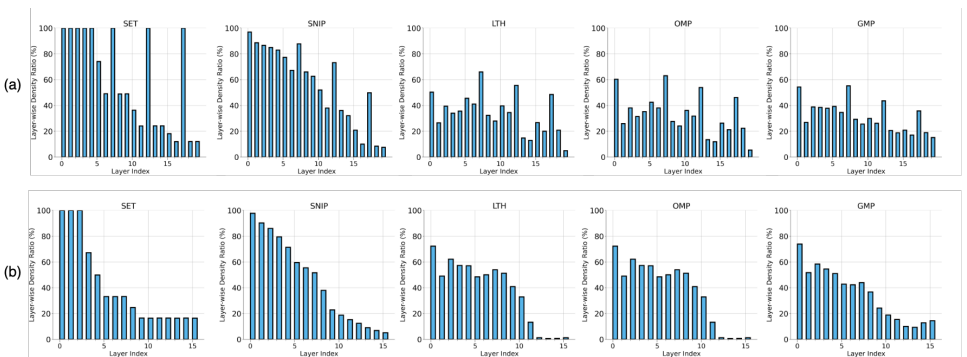


Figure 7: Layer-wise density ratio comparison. This figure compares the layer-wise density ratios of different SNN methods applied to ResNet18 on CIFAR-100 with high EL2N score samples (a), and VGG19 on CIFAR-100 under common corruption scenarios (b) at an overall sparsity ratio of 0.8 and a data ratio of 0.3.

We visualize the layer-wise density ratios for SNN models, focusing on the convolutional layers of CNN architectures. Figure 7 (a) and (b) show notable variations in layer-wise




Figure 8: Accuracy comparison of SNN variants on ResNet18 on CIFAR-100 with high EL2N scores (a) and corruption samples (b), at 0.8 sparsity ratio and 0.3 data ratio. Comparative analysis of training FLOPs, test accuracy, and parameters for SNNs vs. dense models (c, d). ResNet18 on CIFAR-100 with EL2N scores (c) and corruption samples (d). Larger circles indicate models with more parameters.

density ratios for ResNet18 when different SNN methods are applied across various sample difficulties. It is worth noting that for SET method, we utilize the Erdős-Rényi-Kernel (ERK) sparse distribution as in [13, 51], a variation of the Erdős-Rényi (ER) model [57]. Methods like SET and SNIP tend to maintain higher density ratios in shallower layers and lower ratios in deeper layers, consistently across scenarios (see Appendix D for more results). The ERK sparse distribution used in SET for ResNet18 and VGG19 results in higher density ratios in shallow layers (details in Appendix A.1). Other SNN methods determine sparse distributions during training and achieve a more uniform distribution compared to SET and SNIP.

4.2 The Role of Layer-wise Density Ratios in SNN Performance

Can SNN models with similar layer-wise density ratios achieve comparable performance? To answer this, we explore three additional SNN variants: (1) OMP_{ERK}: Maintains the same layer-wise density ratio as the ERK distribution. It involves one-shot pruning on a pre-trained model using weight magnitudes, followed by retraining with a full learning rate schedule. (2) Random Pruning: Adheres to the ERK distribution but uses random sparse topology initialized with binary masks and no pre-trained initialization. The sparse topology remains fixed throughout training. (3) Uniform Pruning: This method contrasts with the ‘Random Pruning’ method by applying a uniform density level across all layers, rather than following the ERK distribution. We trained these methods across different networks and datasets that vary in sample difficulty, allowing for a comprehensive comparison of their performance. The results of these experiments are presented in Figure 8 and Appendix E.

The results reveal that SNNs employing an ERK layer-wise distribution perform similarly to or even better than their dense counterparts. Both ERK and SNIP methods prioritize higher densities in shallower layers, which significantly enhances performance compared to uniform pruning. This finding underscores the role of layer-wise density ratios in improving learning in sparse networks. Strategically allocating higher densities to the shallower layers may help capture essential features in the network, thereby improving the overall performance in sparse settings.

Does the layer-wise density ratio tell the whole story when training with hard samples? The OMP method, which has a lower density in shallower layers, still delivers decent performance (see Figure 4 (a)). This could be attributed to the pre-trained dense model initialization, which has learned to capture essential features in the shallower layers. In contrast,

LTH, which starts from scratch, performs worse than other methods, even with similar density ratios to OMP. This suggests that lower density in the shallower layers can be effective when supported by pre-trained initialization. However, allocating higher densities to the shallower layers is generally better for SNNs, particularly when training from scratch without pre-trained weights.

4.3 SNNs Win Twice When Learned from Hard Examples

In this section, we compare the performance of various SNN methods, considering not only test accuracy but also computational cost. Figure 8 shows the top 5 SNNs for each sparse method, ranked by test accuracy at varying sparsity levels, along with comparisons of training FLOPs and parameters, based on ResNet18 on the CIFAR-100 dataset (at 0.3 data ratio). LTH and OMP methods, which derive sparse topology from pre-trained dense models, require more training FLOPs but with fewer parameters can match dense models in test accuracy at certain sparsity levels. Meanwhile, SET and SNIP, which establish sparse topology before or early in training respectively, outperform dense models in test accuracy while significantly reducing training FLOPs and parameters.

5 Conclusion

This paper provides an empirical analysis of Sparse Neural Networks (SNNs) trained on challenging samples, showing that SNNs often match or exceed the accuracy of dense models at certain sparsity levels while using fewer computational resources in terms of training FLOPs and parameters. This advantage is especially significant in limited data contexts. We find that SNNs with denser connections in shallower layers typically perform better, particularly when training starts from scratch. Future work will focus on exploring additional model efficiency methods like structured pruning.

6 Acknowledgements

This research is part of the research program ‘MegaMind - Measuring, Gathering, Mining and Integrating Data for Self-management in the Edge of the Electricity System’, (partly) financed by the Dutch Research Council (NWO) through the Perspectief program under number P19-25. This research used the Dutch national e-infrastructure with the support of the SURF Cooperative, using grant no. EINF-8980. This work is licensed under Creative Commons Attribution 4.0 International. To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>.

References

- [1] Alessandro Achille, Matteo Rovere, and Stefano Soatto. Critical learning periods in deep networks. In *International Conference on Learning Representations*, 2018.
- [2] Chirag Agarwal, Daniel D’souza, and Sara Hooker. Estimating example difficulty using variance of gradients. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10368–10378, 2022.

- [3] Naveed Akhtar and Ajmal Mian. Threat of adversarial attacks on deep learning in computer vision: A survey. *Ieee Access*, 6:14410–14430, 2018.
- [4] Mike Ashby, Christiaan Baaij, Peter Baldwin, Martijn Bastiaan, Oliver Bunting, Aiken Cairncross, Christopher Chalmers, Liz Corrigan, Sam Davis, N. van Doorn, Jon Fowler, Graham Hazel, Basile Henry, David Page, Jonny Shipton, and S. C. Steenkamp. Exploiting unstructured sparsity on next-generation datacenter hardware. 2019.
- [5] Robert Baldock, Hartmut Maennel, and Behnam Neyshabur. Deep learning through the lens of example difficulty. *Advances in Neural Information Processing Systems*, 34:10876–10889, 2021.
- [6] Anirban Chakraborty, Manaar Alam, Vishal Dey, Anupam Chattopadhyay, and Debdeep Mukhopadhyay. A survey on adversarial attacks and defences. *CAAI Transactions on Intelligence Technology*, 6(1):25–45, 2021.
- [7] Tianlong Chen, Jonathan Frankle, Shiyu Chang, Sijia Liu, Yang Zhang, Zhangyang Wang, and Michael Carbin. The lottery ticket hypothesis for pre-trained bert networks. *Advances in neural information processing systems*, 33:15834–15846, 2020.
- [8] Tianlong Chen, Zhenyu Zhang, Sijia Liu, Shiyu Chang, and Zhangyang Wang. Robust overfitting may be mitigated by properly learned smoothing. In *International Conference on Learning Representations*, 2020.
- [9] Tianlong Chen, Zhenyu Zhang, pengjun wang, Santosh Balachandra, Haoyu Ma, Zehao Wang, and Zhangyang Wang. Sparsity winning twice: Better robust generalization from more efficient training. In *International Conference on Learning Representations*, 2022.
- [10] Yu-Hsin Chen, Tien-Ju Yang, Joel Emer, and Vivienne Sze. Eyeriss v2: A flexible accelerator for emerging deep neural networks on mobile devices. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 9(2):292–308, 2019.
- [11] Selima Curci, Decebal Constantin Mocanu, and Mykola Pechenizkiyi. Truly sparse neural networks at scale. *arXiv preprint arXiv:2102.01732*, 2021.
- [12] Tim Dettmers and Luke Zettlemoyer. Sparse networks from scratch: Faster training without losing performance, 2020.
- [13] Utku Evci, Trevor Gale, Jacob Menick, Pablo Samuel Castro, and Erich Elsen. Rigging the lottery: Making all tickets winners. In *International Conference on Machine Learning*, pages 2943–2952. PMLR, 2020.
- [14] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2019.
- [15] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M Roy, and Michael Carbin. Stabilizing the lottery ticket hypothesis. *arXiv preprint arXiv:1903.01611*, 2019.
- [16] Trevor Gale, Erich Elsen, and Sara Hooker. The state of sparsity in deep neural networks. *arXiv preprint arXiv:1902.09574*, 2019.

- [17] Thomas George, Guillaume Lajoie, and Aristide Baratin. Lazy vs hasty: linearization in deep networks impacts learning schedule based on example difficulty. *arXiv preprint arXiv:2209.09658*, 2022.
- [18] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. *Advances in neural information processing systems*, 28, 2015.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [20] Muyang He, Shuo Yang, Tiejun Huang, and Bo Zhao. Large-scale dataset pruning with dynamic uncertainty. *arXiv preprint arXiv:2306.05175*, 2023.
- [21] Zheng He, Zeke Xie, Quanzhi Zhu, and Zengchang Qin. Sparse double descent: Where network pruning aggravates overfitting. In *International Conference on Machine Learning*, pages 8635–8659. PMLR, 2022.
- [22] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019.
- [23] Pavel Izmailov, Polina Kirichenko, Nate Gruver, and Andrew G Wilson. On feature learning in the presence of spurious correlations. *Advances in Neural Information Processing Systems*, 35:38516–38532, 2022.
- [24] Siddhant Jayakumar, Razvan Pascanu, Jack Rae, Simon Osindero, and Erich Elsen. Top-kast: Top-k always sparse training. *Advances in Neural Information Processing Systems*, 33:20744–20754, 2020.
- [25] Tian Jin, Michael Carbin, Dan Roy, Jonathan Frankle, and Gintare Karolina Dziugaite. Pruning’s effect on generalization through the lens of training and regularization. *Advances in Neural Information Processing Systems*, 35:37947–37961, 2022.
- [26] Angelos Katharopoulos and François Fleuret. Not all samples are created equal: Deep learning with importance sampling. In *International conference on machine learning*, pages 2525–2534. PMLR, 2018.
- [27] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [28] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.
- [29] Namhoon Lee, Thalaisyasingam Ajanthan, and Philip Torr. SNIP: SINGLE-SHOT NETWORK PRUNING BASED ON CONNECTION SENSITIVITY. In *International Conference on Learning Representations*, 2019.
- [30] Shiwei Liu, Decebal Constantin Mocanu, Amarsagar Reddy Ramapuram Matavalam, Yulong Pei, and Mykola Pechenizkiy. Sparse evolutionary deep learning with over one million artificial neurons on commodity hardware. *Neural Computing and Applications*, 33:2589–2604, 2021.

- [31] Shiwei Liu, Lu Yin, Decebal Constantin Mocanu, and Mykola Pechenizkiy. Do we actually need dense over-parameterization? in-time over-parameterization in sparse training. In *International Conference on Machine Learning*, pages 6989–7000. PMLR, 2021.
- [32] Shiwei Liu, Tianlong Chen, Zhenyu Zhang, Xuxi Chen, Tianjin Huang, AJAY KUMAR JAISWAL, and Zhangyang Wang. Sparsity may cry: Let us fail (current) sparse neural networks together! In *The Eleventh International Conference on Learning Representations*, 2023.
- [33] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.
- [34] Sören Mindermann, Jan M Brauner, Muhammed T Razzak, Mrinank Sharma, Andreas Kirsch, Winnie Xu, Benedikt Höltingen, Aidan N Gomez, Adrien Morisot, Sebastian Farquhar, et al. Prioritized training on points that are learnable, worth learning, and not yet learnt. In *International Conference on Machine Learning*, pages 15630–15649. PMLR, 2022.
- [35] Baharan Mirzasoleiman, Kaidi Cao, and Jure Leskovec. Coresets for robust training of deep neural networks against noisy labels. *Advances in Neural Information Processing Systems*, 33:11465–11477, 2020.
- [36] Decebal Constantin Mocanu, Elena Mocanu, Phuong H Nguyen, Madeleine Gibescu, and Antonio Liotta. A topological insight into restricted boltzmann machines. *Machine Learning*, 104:243–270, 2016.
- [37] Decebal Constantin Mocanu, Elena Mocanu, Peter Stone, Phuong H Nguyen, Madeleine Gibescu, and Antonio Liotta. Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. *Nature communications*, 9(1):2383, 2018.
- [38] Hesham Mostafa and Xin Wang. Parameter efficient training of deep convolutional neural networks by dynamic sparse reparameterization. In *International Conference on Machine Learning*, pages 4646–4655. PMLR, 2019.
- [39] Niklas Muennighoff, Alexander M Rush, Boaz Barak, Teven Le Scao, Nouamane Tazi, Aleksandra Piktus, Sampo Pyysalo, Thomas Wolf, and Colin Raffel. Scaling data-constrained language models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [40] Seyed Majid Naji, Azra Abtahi, and Farokh Marvasti. Efficient sparse artificial neural networks. *arXiv preprint arXiv:2103.07674*, 2021.
- [41] Ozan Özdenizci and Robert Legenstein. Training adversarially robust sparse networks via bayesian connectivity sampling. In *International Conference on Machine Learning*, pages 8314–8324. PMLR, 2021.
- [42] Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. Deep learning on a data diet: Finding important examples early in training. *Advances in Neural Information Processing Systems*, 34:20596–20607, 2021.

- [43] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [44] Sylvestre-Alvise Rebuffi, Sven Gowal, Dan Andrei Calian, Florian Stimberg, Olivia Wiles, and Timothy A Mann. Data augmentation can improve robustness. *Advances in Neural Information Processing Systems*, 34:29935–29948, 2021.
- [45] Alex Renda, Jonathan Frankle, and Michael Carbin. Comparing rewinding and fine-tuning in neural network pruning. In *International Conference on Learning Representations*, 2020.
- [46] Leslie Rice, Eric Wong, and Zico Kolter. Overfitting in adversarially robust deep learning. In *International Conference on Machine Learning*, pages 8093–8104. PMLR, 2020.
- [47] Joshua David Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. Contrastive learning with hard negative samples. In *International Conference on Learning Representations*, 2021.
- [48] Shiori Sagawa, Aditi Raghunathan, Pang Wei Koh, and Percy Liang. An investigation of why overparameterization exacerbates spurious correlations. In *International Conference on Machine Learning*, pages 8346–8356. PMLR, 2020.
- [49] Subin Sahayam, John Zakkam, and Umarani Jayaraman. Can we learn better with hard samples? *CoRR*, abs/2304.03486, 2023.
- [50] Connor Shorten and Taghi M Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of big data*, 6(1):1–48, 2019.
- [51] Sidak Pal Singh and Dan Alistarh. Woodfisher: Efficient second-order approximation for neural network compression. *Advances in Neural Information Processing Systems*, 33:18098–18109, 2020.
- [52] Vasu Singla, Sahil Singla, Soheil Feizi, and David Jacobs. Low curvature activations reduce overfitting in adversarial training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16423–16433, 2021.
- [53] David Stutz, Matthias Hein, and Bernt Schiele. Relating adversarially robust generalization to flat minima. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7807–7817, 2021.
- [54] Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A Smith, and Yejin Choi. Dataset cartography: Mapping and diagnosing datasets with training dynamics. *arXiv preprint arXiv:2009.10795*, 2020.
- [55] Hidenori Tanaka, Daniel Kunin, Daniel L Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. *Advances in neural information processing systems*, 33:6377–6389, 2020.

- [56] Mariya Toneva, Alessandro Sordoni, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J. Gordon. An empirical study of example forgetting during deep neural network learning. In *International Conference on Learning Representations*, 2019.
- [57] Mukund Varma T, Xuxi Chen, Zhenyu Zhang, Tianlong Chen, Subhashini Venugopalan, and Zhangyang Wang. Sparse winning tickets are data-efficient image recognizers. *Advances in Neural Information Processing Systems*, 35:4652–4666, 2022.
- [58] Chaoqi Wang, Roger Grosse, Sanja Fidler, and Guodong Zhang. Eigendamage: Structured pruning in the kronecker-factored eigenbasis. In *International conference on machine learning*, pages 6566–6575. PMLR, 2019.
- [59] Chaoqi Wang, Guodong Zhang, and Roger Grosse. Picking winning tickets before training by preserving gradient flow. In *International Conference on Learning Representations*, 2020.
- [60] Peiqi Wang, Yu Ji, Chi Hong, Yongqiang Lyu, Dongsheng Wang, and Yuan Xie. Snram: An efficient sparse neural network computation architecture based on resistive random-access memory. In *2018 55th ACM/ESDA/IEEE Design Automation Conference (DAC)*, pages 1–6, 2018.
- [61] Boqian Wu, Qiao Xiao, Shiwei Liu, Lu Yin, Mykola Pechenizkiy, Decebal Constantin Mocanu, Maurice Van Keulen, and Elena Mocanu. E2enet: Dynamic sparse feature fusion for accurate and efficient 3d medical image segmentation. *arXiv preprint arXiv:2312.04727*, 2023.
- [62] Qiao Xiao, Boqian Wu, Yu Zhang, Shiwei Liu, Mykola Pechenizkiy, Elena Mocanu, and Decebal Constantin Mocanu. Dynamic sparse network for time series classification: Learning what to “see”. *Advances in Neural Information Processing Systems*, 35: 16849–16862, 2022.
- [63] Qiao Xiao, Pingchuan Ma, Adriana Fernandez-Lopez, Boqian Wu, Lu Yin, Stavros Petridis, Mykola Pechenizkiy, Maja Pantic, Decebal Constantin Mocanu, and Shiwei Liu. Dynamic data pruning for automatic speech recognition. *arXiv preprint arXiv:2406.18373*, 2024.
- [64] Geng Yuan, Xiaolong Ma, Wei Niu, Zhengang Li, Zhenglun Kong, Ning Liu, Yifan Gong, Zheng Zhan, Chaoyang He, Qing Jin, et al. Mest: Accurate and fast memory-economic sparse training framework on the edge. *Advances in Neural Information Processing Systems*, 34:20838–20850, 2021.
- [65] Wenyuan Zeng and Raquel Urtasun. Mlprune: Multi-layer pruning for automated neural network compression. 2018.
- [66] Aojun Zhou, Yukun Ma, Junnan Zhu, Jianbo Liu, Zhijie Zhang, Kun Yuan, Wenxiu Sun, and Hongsheng Li. Learning n:m fine-grained structured sparse neural networks from scratch. In *International Conference on Learning Representations*, 2021.
- [67] Michael Zhu and Suyog Gupta. To prune, or not to prune: exploring the efficacy of pruning for model compression. *arXiv preprint arXiv:1710.01878*, 2017.