

DSD²: Can We Dodge Sparse Double Descent and Compress the Neural Network Worry-Free?

Victor Quéto, Enzo Tartaglione

LTCI, Télécom Paris, Institut Polytechnique de Paris, France
 {name.surname}@telecom-paris.fr

Abstract

Neoteric works have shown that modern deep learning models can exhibit a sparse double descent phenomenon. Indeed, as the sparsity of the model increases, the test performance first worsens since the model is overfitting the training data; then, the overfitting reduces, leading to an improvement in performance, and finally, the model begins to forget critical information, resulting in underfitting. Such a behavior prevents using traditional early stop criteria.

In this work, we have three key contributions. First, we propose a learning framework that avoids such a phenomenon and improves generalization. Second, we introduce an entropy measure providing more insights into the insurgence of this phenomenon and enabling the use of traditional stop criteria. Third, we provide a comprehensive quantitative analysis of contingent factors such as re-initialization methods, model width and depth, and dataset noise. The contributions are supported by empirical evidence in typical setups. Our code is available at <https://github.com/VGCQ/DSD2>.

Introduction

Nowadays, deep neural networks are one of the most employed algorithms when required to solve complex tasks. In particular, their generalization capability allowed them to establish new state-of-the-art performance in domains like computer vision (He et al. 2016; Dosovitskiy et al. 2021) and natural language processing (Vaswani et al. 2017; Brown et al. 2020), showing as well promising capability in very complex hybrid tasks, like text-to-image generation (Ramesh et al. 2022; Saharia et al. 2022). The problem of optimally sizing these models is relevant to the vastly distributed employment of deep neural networks on edge devices (Chen and Ran 2019; Lin et al. 2022), posing questions about power consumption and hardware complexity (Goel et al. 2020; Luo et al. 2022).

It was general knowledge that the more a model is over-parametrized, the easier it will overfit the training set, entering the *memorization phase*: the model memorizes the single samples in the training set, learning also a wrong set of features. Such a phenomenon harms the model’s generalization, worsening its performance on unseen data (Liu

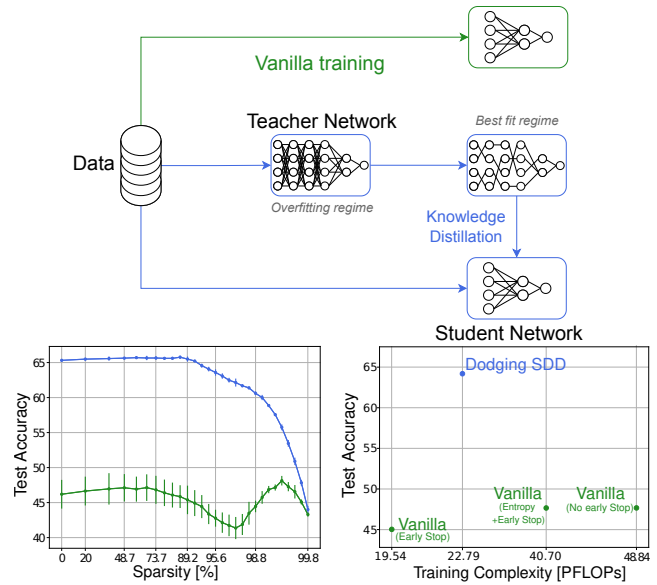


Figure 1: Distilling knowledge from a sparse teacher grants access to solutions (for the student model) where SDD is dodged, also saving computation.

et al. 2020). Recently, a new surprising phenomenon, *double descent* (DD) (Belkin et al. 2019), has been observed for extremely over-parametrized models: beyond the traditional over-fitting regime, while continuing to increase the size of the model, the generalization gap between train and test performance inverts trend, and narrows the more, the larger the model is (Nakkiran et al. 2020).

The DD phenomenon raises questions about how to optimally size the model to have the best performance (while having the minimum size). Many approaches have indeed proposed the use of regularization functions to relieve DD in models for regression and classification tasks. However, when moving to real applications, the complexity of optimally tuning the regularization hyper-parameters and learning optimal early-stop discourages their use (Kan, Nagy, and Ruthotto 2020). While DD’s typical analysis concerns moving from small models to big ones, a recent work observed a similar phenomenon also moving from an over-parametrized model backward to a smaller one (He et al.

2022). That is possible thanks to pruning, which iteratively removes parameters from the model. Intuitively, while pruning removes parameters from the model, while at first the performance is enhanced, it will enter a second phase where its inevitable worsening is met (Han et al. 2015; Quéu and Tartaglione 2023). Such effect took the name of *sparse double descent* (SDD): is it inevitable? Can we properly regularize the model toward test performance enhancement? What is the underlying explanation of the SDD phenomenon?

We summarize our contributions as follows.

- To the best of our knowledge, we propose the first approach avoiding SDD, and providing a model with a good performance in terms of validation/test accuracy, consistently. More specifically, we regularize a student model distilling knowledge from a sparse teacher (in its best validation accuracy region), observing that the student dodges SDD (Fig. 1). Interestingly, we observe that this happens even when distilling knowledge from a non-pruned teacher.
- We study SDD from the perspective of “neuron’s states”: we calculate the entropy of the activations in the model, observing a correlation between the *interpolation regime* (where the SDD occurs) and the entropy’s flatness. When leaving such a region, thus entering the *classical regime* where the traditional bias/variance trade-off occurs, the entropy monotonically decreases: a check on this measure enables-back the use of early stop criteria, which saves training computational cost (rightmost Fig. 1).
- We propose a quantitative study on some open questions, in particular: i) is DD/SDD still occurring when models increase in depth? ii) do we find the best validation/test performance configuration in models extremely over-parametrized or right before under-fitting? iii) is there a big difference between setting back parameters to their initial value (rewinding), randomly initializing, and not perturbing the model’s parameters after pruning?

Related Works

The Real World Is Noisy In the real world, data acquisition is often noisy (Gupta and Gupta 2019), stemming from data collection or labeling. Concerning the annotation noise, many works proposed solutions to prevent the learning of wrong feature sets: for example, (Li et al. 2017) propose a unified distillation framework by leveraging the knowledge learned from a small clean dataset and semantic knowledge graph to correct the noisy labels. Other works find solutions inspired by the benefits of noise in the nervous system: (Arani, Sarfraz, and Zonooz 2021), for example, show that injecting constructive noise at different levels in the collaborative learning framework enables training the model effectively and distills desirable characteristics in the student model. More specifically, they propose methods to minimize the performance gap between a compact and a large model, to train high-performance compact adversarially-robust models.

Since a single image may belong to several categories, different samples can suffer from varying intensities of label noise. (Xu et al. 2020) proposed a simple yet effective feature normalized knowledge distillation that introduces

the sample-specific correction factor to replace the temperature. (Kaiser et al. 2022) developed a teacher-student approach that identifies the tipping point between good generalization and overfits, thus estimating the noise in the training data with Otsu’s algorithm. Other works focus more on the label’s prediction robustness: (Sau and Balasubramanian 2016), for example, introduced a simple method that helps the student to learn better and produces results closer to the teacher network by injecting noise and perturbing the logit outputs of the teacher. With this setup, the noise simulates a multi-teacher setting and produces the effect of a regularizer due to its presence in the loss layer. To simulate noise in the labels, a typical approach is to manually inject noise in some well-annotated, standard datasets like MNIST and CIFAR-10/100 (Nakkiran et al. 2020; He et al. 2022). AI security works use a similar setup as well, where noise is injected parametrically to analyze the model’s robustness against attacks. In a nutshell, these attacks propose adversarial representations of the data and check the model’s performance - so this setup is named “adversarial learning” (Miller, Xiang, and Kesidis 2020).

Double Descent in Classification Tasks Considering the presence of labeling noise, the occurrence of DD is a real threat. DD has already been reported in various machine learning models, like decision trees, random features (Meng, Yao, and Cao 2022), linear regression (Muthukumar et al. 2020)(Belkin, Hsu, and Xu 2020; Hastie et al. 2022), and deep neural networks (Yilmaz and Heckel 2022). For classification tasks, the test error of standard deep networks, like the ResNet architecture, trained on image classification datasets, consistently follows a double descent curve both when label noise is injected (CIFAR-10), and in some cases, even without any label noise injection (CIFAR-100) (Yilmaz and Heckel 2022). (Nakkiran et al. 2020) show that double descent occurs not just as a function of model size when increasing the model width, but also as a function of the number of training epochs. The double descent phenomenon has been extensively studied under the spectrum of over-parametrization (Nakkiran et al. 2020; Chang et al. 2021). (He et al. 2022) observe the occurrence of double descent not only in the traditional setups but also when unstructurally pruning a dense model, observing the SDD phenomenon. Working on a related research question, and motivating the importance of further studying the SDD, (Chang et al. 2021) addressed the important question of whether it could be more convenient to train a small model directly, or quite first train a larger one and then prune it. In this work, the authors provide convincing evidence that the latter strategy is winning, toward enhanced model performance in sparsified regimes. Cotter et al. (2021), as a contrast to this work’s purpose, exploited the DD phenomenon in a self-supervised setup, to assign pseudo-labels to a large held-out dataset. While this work sought to exploit DD, our goal is different: our objective is to avoid the sparse double descent, towards enhanced performance on the final student model, on the same task and dataset as the teacher. In the next section, we present and show the occurrence of SDD.




Figure 2: Performance of ResNet-18 with different amount of noise ε on CIFAR-10 (a) and CIFAR-100 (b). I: Light Phase. II: Critical Phase. III: Sweet Phase. IV: Collapsed Phase.

Model Size and Sparse Double Descent

Background on Neural Network’s Pruning

Neural network pruning aims to reduce a large network while maintaining accuracy by removing irrelevant weights, filters, or other structures, from neural networks. All the pruning algorithms removing weights without explicitly considering the neural network’s structure are typically named *unstructured* pruning methods. Various unstructured pruning methods exist and can be divided into magnitude-based, where the ranking for the parameters to prune is based on their magnitude (Han et al. 2015; Louizos, Welling, and Kingma 2018; Zhu and Gupta 2018), and gradient-based, where the ranking or the penalty term is a function of the gradient magnitude (or to higher order derivatives) (Lee, Ajanthan, and Torr 2019; Tartaglione et al. 2022). The general comparison between the effectiveness of any of the reported approaches is reported by (Blalock et al. 2020) and, although complex pruning approaches exist, the simple magnitude-based one, in general, is considered a good trade-off between complexity and competitiveness (Gale, Elsen, and Hooker 2019): hence, we will focus on this one.¹

Training first an over-parametrized model, and then pruning it, leads to an improved generalization. In particular, (Chang et al. 2021) analyzes the beneficial effects of pruning, and then compares the performance achieved by pruned

¹(He et al. 2022) showed that magnitude, gradient-based, and random pruning achieve similar performance for the same setup as we consider in this work.

models to shallow vanilla ones. This work motivates the quest for investigating SDD, looking for the highest possible performance on unseen data.

Pruning Exhibits Sparse Double Descent

Setup The trained model \mathcal{M} on the train set $\mathcal{D}_{\text{train}}$ (whose performance is evaluated on the validation set \mathcal{D}_{val}) consists of L layers, having $\mathbf{w}^{\mathcal{M}}$ as its set of parameters, and $\mathbf{w}_l^{\mathcal{M}}$ indicates those belonging to the l -th layer. When we prune the ζ -th fraction of parameters from the model, the parameters are projected to a parameter sub-space, according to a threshold on the quantile function $\mathcal{Q}_{\mathcal{M}}(\cdot)$, computed on the absolute values for parameters in \mathcal{M} .

The overall approach employed to reduce the dimensionality of the trained \mathcal{M} follows these steps. The first step is to train the dense model. Until it has reached the desired sparsity percentage ζ_{wall} , the model is iteratively pruned using some pruning strategy (i.e. magnitude pruning, following (He et al. 2022)), perturbed (weights can be rewound to initialization, randomly re-initialized, or not perturbed at all), and the sparse model is re-trained on $\mathcal{D}_{\text{train}}$. When ζ_{wall} is reached, the model parameters \mathbf{w}_{best} , which achieve the best performance on the validation set \mathcal{D}_{val} are returned.

Experiments As in (He et al. 2022), SDD is, consistently, found in our experiments. In particular, Fig. 2a and Fig. 2b show double descent of a ResNet-18 model trained on CIFAR-10 and CIFAR-100, respectively. The four phases reported by (He et al. 2022) are displayed in Fig. 2 (first a




Figure 3: Performance of ResNet-18 on CIFAR-100 with $\varepsilon = 10\%$ when retrained from either the original initialization (lottery ticket), a random re-initialization, or from the last configuration achieved before pruning.

light phase, secondly a *critical* phase, thirdly the *sweet* phase and finally, the *collapsed* phase).

Better Low Parametrization or Extreme Over-parametrization?

Currently, there is a big debate about whether some simple techniques, like early stopping, are sufficient to achieve great generalization performance. In particular, (Rice, Wong, and Kolter 2020) studies the extreme over-parametrization for adversarially trained deep networks. The authors observe that overfitting the training set harms robust performance to a large degree in adversarially robust training across multiple datasets: this can be fought by simply using early stopping. To mitigate the robust overfitting, (Chen et al. 2021) propose self-training to smoothen the logits, combined with stochastic weight averaging trained by the same model, the other performing stochastic weight averaging (Izmailov et al. 2018). Although these works focus on DD, similar effects can be drawn to SDD: is it always true that the best model is in the sweet phase?

In our results presented in Fig. 2, we observe a correlation between the best model (marked with \star) and the noise in the dataset: on CIFAR-10, w_{best} is located in the sweet phase for $\varepsilon \in \{20\%, 50\%\}$, while on CIFAR-100, w_{best} is in the sweet phase only for $\varepsilon = 50\%$. Therefore, we empirically observe a correlation between the amount of noise in the training set and the location of w_{best} : for small noise (i.e. $< 15\%$), w_{best} is located in the Light Phase (I). As ε exceeds 15% , w_{best} is consistently found in the Sweet Phase (III). A further analysis of this shift is conducted in Appendix.

Critical Phase Occurrence

(Frankle and Carbin 2019) conjectured, under the so-called *lottery ticket hypothesis*, that a large network contains smaller sub-networks, which can be trained in isolation without performance loss. Their approach consists of

rewinding the parameters to their original value every pruning iteration and is one of the most popular for neural network pruning (Malach et al. 2020; Zhang et al. 2021).

We compare, in Fig. 3, different strategies to introduce perturbations in the model’s parameters - random re-initialization, rewinding (in the lottery ticket hypothesis fashion), and not introducing any perturbation-, averaged on three seeds. We observe, in all the above-mentioned cases, the rise of the sparse double descent. However, we notice a gap in the performance between the different perturbation approaches. While randomly re-initializing the model leads to worse results, rewinding and not introducing any perturbation exhibit slightly different behavior, depending on the pruning regime. In particular, rewinding can marginally achieve a better performance, but the performance decays faster than not-perturbing. Moreover, we investigate whether it could be more convenient to let the model remain in the neighborhood of the same local minimum found by the dense model. Towards this end, we propose an experiment where we do not perturb the learned parameters, and we scale down the learning rate. In this scenario, we observe the performance consistently deteriorated. Our finding suggests that with rewinding we are post-posing, in the ζ plane, the critical phase, and without introducing any perturbation we are further post-posing it: we will use a “no perturbation” scheme for all our experiments.

An Entropy-Based Interpretation to the Sparse Double Descent

Analyzing the SDD phenomenon from a learning dynamics perspective raises questions related to the so-called *information bottleneck theory* (Tishby et al. 1999; Tishby and Zaslavsky 2015). In particular, this approach estimates the mutual information between the information processed by the layers and the input and output variables. Hence, it is possible to calculate optimal theoretical limits and set the bars for the generalization error. Several works have followed this theory, observing differences in the learning dynamics for different activation functions employed (Saxe et al. 2018), improving the estimation approach (Pan et al. 2021), and verifying that such theory admits DD in regression tasks (Ngampruetikorn and Schwab 2022). Inspired by these works, we formulate the following observation.

Observation 1 *As the size of the model, in terms of the number of parameters per neuron, shrinks from the light phase, the entropy of the features inside the model is stationary (as small adjustments in the parameters are needed). As we leave the interpolation regime right after the interpolation threshold, the entropy begins to decrease.*

To empirically verify this observation, we can visualize the entropy of the activations in the model. Considering that our observation is not constrained to any validation/test set, we will perform the measures directly on the training set. We define the average neuron’s entropy in the l -th layer as

$$\bar{\mathcal{H}}_{l|\mathcal{D}_{\text{train}}} = -\frac{1}{N_l} \sum_{i=1}^{N_l} \sum_{\zeta \in \{0,1\}} p(s_{l,i}^{\zeta}) \log \left[p(s_{l,i}^{\zeta}) \right], \quad (1)$$

where $p(s_{l,i}^\zeta)$ is the probability (in a frequentism sense, over $\mathcal{D}_{\text{train}}$) the i -th neuron (ReLU-activated) in the l -th layer to be in the negative region ($\zeta = 0$) or in the positive one ($\zeta = 1$), similarly to how done in (Liao et al. 2023; Spadaro et al. 2023). Fig. 2 displays the variation of the entropy, averaged across all the model’s layers. For the considered examples, we corroborate our observation, enabling-back the use of early-stop criteria jointly with the entropy. Indeed, the entropy stays stationary and then decreases when the model enters the classical regime (entering the sweet phase). Using traditional early criteria starting from this regime saves training computation as the pruning/training process is stopped when the performance decreases.

Generalization Gap in Deep Double Descent: Relationships between DD and SDD

In this section, we analyze the occurrence of the sparse double descent concerning the model’s size. To carry out the following experiments, we have defined a multi-configurable “VGG-like” model: we can generate multiple architectures depending on the depth δ (the larger, the deeper the model) and the number of convolutional filters per layer 2^γ (the larger, the wider the model).

Results Fig. 4a shows the results at growing depth, with fixed $\gamma = 5$. The sparse double descent becomes more and more evident for growing depths, while for shallow models we can only observe the traditional regime of sweet and collapsed phase. We observe a similar trend also in Fig. 4b, where the depth of the model is set to $\delta = 1$ (two convolutional layers and one fully connected) while the width of the layer is increasing. With a fixed depth, increasing the width of the layers also reveals a double descent phenomenon: with low width, no SDD is observed, but it can be observed as γ is increased.

Typical analyses are performed, in the literature, in terms of the number of samples in the training set, or in terms of the width of the layers (Nakkiran et al. 2020; Chang et al. 2021; Chen, Wang, and Kyrillidis 2021). As the phenomenon also rises as a function of the model’s depth, we believe SDD is more related to the number of parameters in a model rather than to the layer’s organization and structure. This could motivate the occurrence of the sparse double descent, as the number of parameters is the varying quantity, while the width of the layer might not change.

Distilling Knowledge to Avoid the Sparse Double Descent

Although a standard ℓ_2 regularization approach can already positively contribute to dodging SDD, it also presents some flaws. (Nakkiran et al. 2021) showed that, for certain linear regression models with isotropic data distribution, optimally-tuned ℓ_2 regularization can achieve monotonic test performance as either the sample size or the model size is grown. However, (Quétu and Tartaglione 2023) highlighted that in some image classification setups, like ResNet-18 on CIFAR datasets, SDD is still noticeable even if ℓ_2 regularization is used. Moreover, this regularization has the

big drawback of sacrificing the performance/sparsity trade-off (Quétu, Milovanović, and Tartaglione 2023). Thus, the need to design a custom regularization to prevent SDD, while maintaining good performance, becomes evident. In this section, we present a learning scheme in which a student model is regularized by distilling knowledge from a sparse teacher in its best validation accuracy region (or even with a dense teacher), observing that the student is dodging SDD.

Background on Knowledge Distillation

Knowledge distillation (KD) transfers the knowledge from a (deep) teacher to a (shallow) student network. There are several knowledge types: response-based, feature-based, and relation-based. While response-based knowledge uses the output of the last layer of the teacher model as the knowledge (Hinton et al. 2015), the output of intermediate layers, which in convolutional neural networks are the feature maps, is used to supervise the training of the student model for feature-based knowledge (Romero et al. 2014). Relation-based knowledge further explores the relationships between different layers or data samples (Yim et al. 2017). In this work, we will focus on the more flexible and easy-to-deploy response-based knowledge.

Distilling the knowledge from a teacher model, besides enhancing the student’s performance, can also implicitly transmit good regularization properties. While (Yuan et al. 2020) observe that good students could also improve the teacher’s performance, and (Furlanello et al. 2018) observe improvement in distilling knowledge from students in a “born again” fashion, (Saglietti and Zdeborová 2022) introduce a formal statistical physics framework that allows an analytic characterization of the properties of KD in shallow neural networks. By using Gaussian Mixture Models, it is shown that, without any fine-tuning at the level of the student loss function, KD allows a transfer of the (possibly fine-tuned) regularization properties of the teacher, even if the two models are mismatched, and even if the regularization strategy in the teacher training is not explicitly known. This work motivates us, in the quest for transmitting an implicit teacher’s regularization (as being either in the well-generalizing light phase or the sweet phase for high label noise) to a student, in the attempt to escape from the SDD. Indeed, (Saglietti and Zdeborová 2022) highlight a noticeable enhancement in the distillation test error when the transferred outputs are produced by an optimal teacher, with respect to a direct ℓ_2 regularization: the gap with the optimal generalization bound is nearly closed. KD enables the student model to reach a performance unachievable with other regularization approaches, inheriting the teacher’s regularization implicitly.

Combining Knowledge Distillation and Network Pruning

Several works have already explored the combination of KD and pruning. Towards generalization enhancement, in the context of working with reduced size of the train set, (Zhou et al. 2022) propose a “progressive feature distribution distillation”, which consists of first collecting a student network

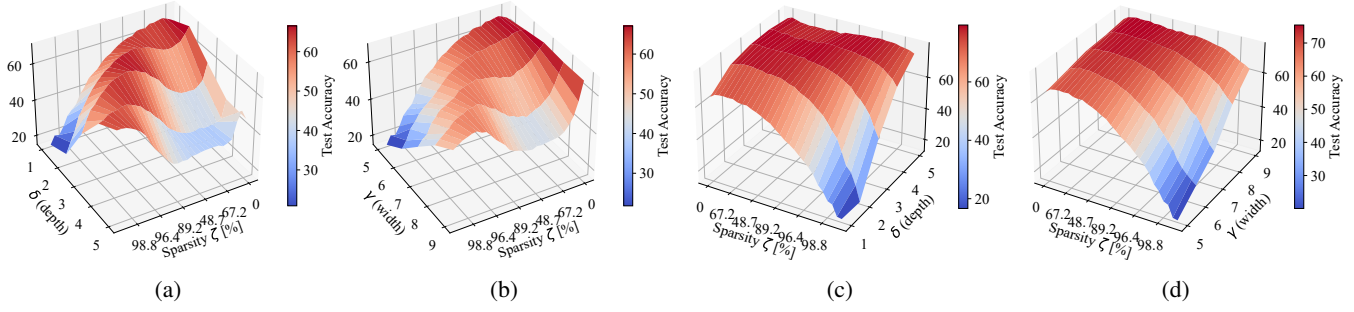


Figure 4: Performance of VGG-like model, vanilla-trained (a, b), distilled from a sparse teacher (c, d), varying the depth δ (a, c) and the width γ (b, d) on CIFAR-10 with $\varepsilon = 50\%$.

by pruning a trained network, and then distilling the feature distribution into the student.

A large slice of work combining the two techniques targets the final model’s size reduction. For example, (Cui, Li, and Zhang 2021) combine structured pruning and dense knowledge distillation techniques to significantly compress an original large language model into a deep compressed shallow network, and (Wei, Hao, and Zhang 2022) follow a similar strategy for speech enhancement. Still working on language models, (Kim and Rush 2016) reveal that standard KD applied to word-level prediction can be effective for Neural Machine Translation, and applied weight pruning on top of it to decrease the number of parameters. Image processing-related works are, more applicative and oriented towards the enhancement of the performance on a sparsified model: (Chen et al. 2022) and (Aghli and Ribeiro 2021) are two works in this context. Leveraging on the beneficial, well-generalizing properties of both pruning and KD, (Park and No 2022) provide several applicative examples where the “prune, then distill” pattern is practically very effective.

Motivated by (Saglietti and Zdeborová 2022) and building on top of (Park and No 2022) (although in a very different context), we formulate the following observation, which will drive our approach towards SDD dodging.

Observation 2 *In an adversarial learning setup with knowledge distillation, the teacher’s response will act as a regularizer for the student model, which makes it avoid the sparse double descent. The student’s performance, in a high-noise setup, is further enhanced when the knowledge is distilled from a sparse, well-performing teacher.*

Approach In general, the objective function to be minimized while training a smaller student network in the KD framework is a linear combination of two losses, in the image classification setup: a standard cross-entropy loss \mathcal{L}_{CE} , which uses “hard” ground truth targets, and the Kullback-Leibler divergence loss \mathcal{L}_{KL} , calculated between the teacher’s and the student’s predictions, eventually scaled by a temperature τ :

$$\mathcal{L} = (1 - \alpha)\mathcal{L}_{CE}(\mathbf{y}^s, \hat{\mathbf{y}}) + \alpha\mathcal{L}_{KL}(\mathbf{y}^s, \mathbf{y}^t, \tau), \quad (2)$$

where the apex “s” and “t” refer to the student and the teacher respectively, $\hat{\mathbf{y}}$ is the ground truth label, and α is

the distillation hyper-parameter impacting on the weighted average between the distillation loss and the student loss. We employ the KD formulation loss as in (Hinton et al. 2015; Kim et al. 2021). To distill the knowledge to some students, we require the teacher model to be already in its best-fit regime, and using the same sparsification process - but changing the objective function to (2)-, we can train and sparsify the student model.

Results For all the presented experiments here, $\alpha = 0.8$ and $\tau = 10$. Fig. 5 shows the results on CIFAR-10 and CIFAR-100 with $\varepsilon \in \{10\%, 20\%, 50\%\}$ using a VGG-like model with $2^\gamma = 32$ and $\delta = 5$ as student model, while ResNet-18 as teacher (whose training and performance is consistent to Fig. 2). We observe, consistently for all the investigated noise rates, that the student model, trained with a vanilla training setup, always reveals the sparse double descent phenomenon. Nevertheless, the same student model, trained using a KD setup, can consistently avoid SDD. Furthermore, even if the knowledge is distilled from a dense teacher, SDD is always avoided, either in the case of best performance in the light or the sweet phase. Observing the entropy for the model, it stays stationary and then decreases when the model enters the classical regime. Hence, once the entropy decreases, if the performance drops, the training can be stopped as the model enters the under-fitting regime.

Computation/performance Trade-off We compare traditional approaches with our proposed scheme on CIFAR-10 with $\varepsilon = 20\%$ in Table 1 in terms of performance achieved and training computational cost for the student model.

In the vanilla case, early stop results in a model achieving low performance and with low sparsity. In order to extract a better performance, apparently, there is no other choice than pruning the model until all of its parameters are completely removed, which costs more than 48 PFLOPs. Our method achieves a model with high sparsity achieving more than 10% improvement in performance with approximately 25% computation less. We believe this is a core applicative contribution in real applications, where annotated datasets are small and noisy, and SDD can be easily observed.

Experiments with Varying Depth and Width In Fig. 4c and Fig. 4d we report the distillation results on CIFAR-10

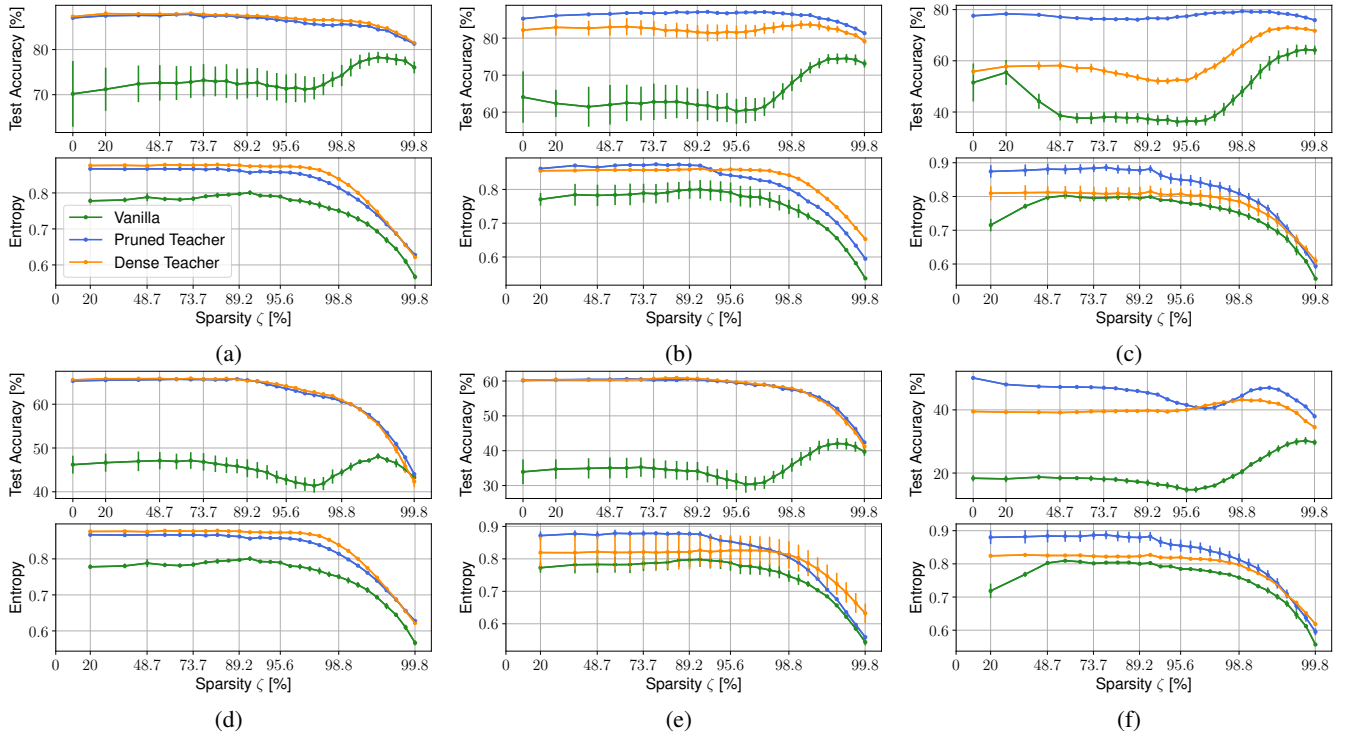


Figure 5: Performance of the VGG-like model on CIFAR-10 (a, b, c) and CIFAR-100 (d, e, f) for different label noises. Left: $\varepsilon = 10\%$. Middle: $\varepsilon = 20\%$. Right: $\varepsilon = 50\%$.

Early stop	Distillation	Distillation from pruned teacher	Training complexity [PFLOPs](↓)	Test accuracy [%](↑)	Sparsity [%](↑)
			48.84	74.52 ± 1.20	99.62
✓			4.88	60.23 ± 3.37	36.00
✓	✓		35.82 (+1.63)	81.52 ± 1.85	99.26
✓	✓	✓	35.82 (+47.21)	86.89 ± 0.16	99.26

Table 1: Performance achieved and training computational cost of the student model with traditional approaches and our proposed scheme on CIFAR-10 with $\varepsilon = 20\%$. The training cost of the teacher model is reported between parenthesis.

for $\varepsilon = 50\%$, with a varying depth and width for the student (results for the other noise values are displayed in Appendix). While in Fig. 4a and Fig. 4b we consistently observe that the vanilla-trained student model exhibits SDD as the width or depth is increased, we persistently notice in Fig. 4c and Fig. 4d, however, that employing KD within the framework, the performance exhibits a monotonic behavior: the sparse double descent is dodged.

Limitations and Future Work Leveraging a knowledge distillation scheme is a successful approach to transmitting good generalization properties from the teacher model to the student and enabling the dodge of SDD on the student itself. However, this method also presents some limits: in resource-constrained schemes, one might not afford to train a large teacher model. Nevertheless, our work is the first to enable the small model to enter the sweet phase instead of the critical phase, which cannot be achieved with traditional regularization strategies. Therefore, we let the exploration of more efficient approaches as future work.

Conclusion

In this paper, we have studied and proposed a solution to dodge the sparse double descent. This phenomenon poses critical questions about where to find the model with the best performance in the low or the high sparsification regime. We observe an interesting correlation between the occurrence of the critical phase and the entropy of the activation’s state for the training set - when not in the classical regime, the entropy is stationary. Although SDD challenges traditional early-stop schemes, the use of the network’s entropy, which indicates the region we are navigating on, enables them back. We further observe that by leveraging a knowledge distillation scheme, not only the good generalization property of the teacher is transmitted to the student, but the student itself is no longer subject to SDD, due to the implicit regularization inherited by the teacher. We hope this work will ignite new research toward the improvement of learning strategies for deep, sparse models.

Acknowledgements

This work was performed using HPC resources from GENCI-IDRIS (Grant 2022-AD011013930). Also, this research was partially funded by Hi!PARIS Center on Data Analytics and Artificial Intelligence.

References

- Aghli, N.; and Ribeiro, E. 2021. Combining weight pruning and knowledge distillation for cnn compression. In *CVPR*.
- Arani, E.; Sarfraz, F.; and Zonooz, B. 2021. Noise as a resource for learning in knowledge distillation. In *WACV*.
- Belkin, M.; Hsu, D.; Ma, S.; and Mandal, S. 2019. Reconciling modern machine-learning practice and the classical bias-variance trade-off. *Proceedings of the National Academy of Sciences*.
- Belkin, M.; Hsu, D.; and Xu, J. 2020. Two models of double descent for weak features. *SIAM Journal on Mathematics of Data Science*.
- Blalock, D.; Gonzalez Ortiz, J. J.; Frankle, J.; and Gutttag, J. 2020. What is the state of neural network pruning? *MLSys*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *NeurIPS*.
- Chang, X.; Li, Y.; Oymak, S.; and Thrampoulidis, C. 2021. Provable benefits of overparameterization in model compression: From double descent to pruning neural networks. In *AAAI*.
- Chen, J.; and Ran, X. 2019. Deep learning with edge computing: A review. *IEEE*.
- Chen, J.; Wang, Q.; and Kyrillidis, A. 2021. Mitigating deep double descent by concatenating inputs. In *CIKM*.
- Chen, L.; Chen, Y.; Xi, J.; and Le, X. 2022. Knowledge from the original network: restore a better pruned network with knowledge distillation. *Complex & Intelligent Systems*.
- Chen, T.; Zhang, Z. A.; Liu, S.; Chang, S.; and Wang, Z. 2021. Robust Overfitting may be mitigated by properly learned smoothening. In *ICLR*.
- Cotter, A.; Menon, A. K.; Narasimhan, H.; Rawat, A. S.; Reddi, S. J.; and Zhou, Y. 2021. Distilling Double Descent. *arXiv preprint arXiv:2102.06849*.
- Cui, B.; Li, Y.; and Zhang, Z. 2021. Joint structured pruning and dense knowledge distillation for efficient transformer model compression. *Neurocomputing*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Hounsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*.
- Frankle, J.; and Carbin, M. 2019. The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks. In *ICLR*.
- Furlanello, T.; Lipton, Z.; Tschannen, M.; Itti, L.; and Anandkumar, A. 2018. Born again neural networks. In *ICML*.
- Gale, T.; Elsen, E.; and Hooker, S. 2019. The state of sparsity in deep neural networks. *arXiv preprint arXiv:1902.09574*.
- Goel, A.; Tung, C.; Lu, Y.-H.; and Thiruvathukal, G. K. 2020. A survey of methods for low-power deep learning and computer vision. In *WF-IoT*.
- Gupta, S.; and Gupta, A. 2019. Dealing with noise problem in machine learning data-sets: A systematic review. *Procedia Computer Science*.
- Han, S.; Pool, J.; Tran, J.; and Dally, W. 2015. Learning both weights and connections for efficient neural network. *NeurIPS*.
- Hastie, T.; Montanari, A.; Rosset, S.; and Tibshirani, R. J. 2022. Surprises in high-dimensional ridgeless least squares interpolation. *Annals of statistics*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- He, Z.; Xie, Z.; Zhu, Q.; and Qin, Z. 2022. Sparse Double Descent: Where Network Pruning Aggravates Overfitting. In *ICML*.
- Hinton, G.; Vinyals, O.; Dean, J.; et al. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Izmailov, P.; Podoprikin, D.; Garipov, T.; Vetrov, D.; and Wilson, A. G. 2018. Averaging weights leads to wider optima and better generalization. *UAI*.
- Kaiser, T.; Ehmann, L.; Reinders, C.; and Rosenhahn, B. 2022. Blind Knowledge Distillation for Robust Image Classification. *arXiv preprint arXiv:2211.11355*.
- Kan, K.; Nagy, J. G.; and Ruthotto, L. 2020. Avoiding The Double Descent Phenomenon of Random Feature Models Using Hybrid Regularization. *arXiv preprint arXiv:2012.06667*.
- Kim, T.; Oh, J.; Kim, N.; Cho, S.; and Yun, S.-Y. 2021. Comparing kullback-leibler divergence and mean squared error loss in knowledge distillation. *IJCAI*.
- Kim, Y.; and Rush, A. M. 2016. Sequence-level knowledge distillation. *EMNLP*.
- Lee, N.; Ajanthan, T.; and Torr, P. 2019. SNIP: Single-shot network pruning based on connection sensitivity. In *ICLR*.
- Li, Y.; Yang, J.; Song, Y.; Cao, L.; Luo, J.; and Li, L.-J. 2017. Learning from noisy labels with distillation. In *ICCV*.
- Liao, Z.; Quéto, V.; Nguyen, V.-T.; and Tartaglione, E. 2023. Can Unstructured Pruning Reduce the Depth in Deep Neural Networks? In *ICCVW*.
- Lin, J.; Zhu, L.; Chen, W.-M.; Wang, W.-C.; Gan, C.; and song han. 2022. On-Device Training Under 256KB Memory. In *NeurIPS*.
- Liu, S.; Niles-Weed, J.; Razavian, N.; and Fernandez-Granda, C. 2020. Early-learning regularization prevents memorization of noisy labels. *NeurIPS*.
- Louizos, C.; Welling, M.; and Kingma, D. P. 2018. Learning Sparse Neural Networks through L_0 Regularization. In *ICLR*.

- Luo, X.; Liu, D.; Kong, H.; Huai, S.; Chen, H.; and Liu, W. 2022. Lightnas: On lightweight and scalable neural architecture search for embedded platforms. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*.
- Malach, E.; Yehudai, G.; Shalev-Schwartz, S.; and Shamir, O. 2020. Proving the lottery ticket hypothesis: Pruning is all you need. In *ICML*.
- Meng, X.; Yao, J.; and Cao, Y. 2022. Multiple Descent in the Multiple Random Feature Model. *arXiv preprint arXiv:2208.09897*.
- Miller, D. J.; Xiang, Z.; and Kesidis, G. 2020. Adversarial Learning Targeting Deep Neural Network Classification: A Comprehensive Review of Defenses Against Attacks. *Proceedings of the IEEE*.
- Muthukumar, V.; Vodrahalli, K.; Subramanian, V.; and Sahai, A. 2020. Harmless interpolation of noisy data in regression. *JSAT*.
- Nakkiran, P.; Kaplun, G.; Bansal, Y.; Yang, T.; Barak, B.; and Sutskever, I. 2020. Deep Double Descent: Where Bigger Models and More Data Hurt. In *ICLR*.
- Nakkiran, P.; Venkat, P.; Kakade, S. M.; and Ma, T. 2021. Optimal Regularization can Mitigate Double Descent. In *ICLR*.
- Ngampruetikorn, V.; and Schwab, D. J. 2022. Information bottleneck theory of high-dimensional regression: relevancy, efficiency and optimality. In *NeurIPS*.
- Pan, Z.; Niu, L.; Zhang, J.; and Zhang, L. 2021. Disentangled information bottleneck. In *AAAI*.
- Park, J.; and No, A. 2022. Prune your model before distill it. In *ECCV*.
- Quétu, V.; Milovanović, M.; and Tartaglione, E. 2023. Sparse Double Descent in Vision Transformers: real or phantom threat? In *ICIAP*.
- Quétu, V.; and Tartaglione, E. 2023. Dodging the Double Descent in Deep Neural Networks. In *ICIP*.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Rice, L.; Wong, E.; and Kolter, J. Z. 2020. Overfitting in adversarially robust deep learning. In *ICML*.
- Romero, A.; Ballas, N.; Kahou, S. E.; Chassang, A.; Gatta, C.; and Bengio, Y. 2014. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*.
- Saglietti, L.; and Zdeborová, L. 2022. Solvable model for inheriting the regularization through knowledge distillation. In *MSML*.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E.; Ghasemipour, S. K. S.; Gontijo-Lopes, R.; Ayan, B. K.; Salimans, T.; Ho, J.; Fleet, D. J.; and Norouzi, M. 2022. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. In *NeurIPS*.
- Sau, B. B.; and Balasubramanian, V. N. 2016. Deep model compression: Distilling knowledge from noisy teachers. *arXiv preprint arXiv:1610.09650*.
- Saxe, A. M.; Bansal, Y.; Dapello, J.; Advani, M.; Kolchinsky, A.; Tracey, B. D.; and Cox, D. D. 2018. On the Information Bottleneck Theory of Deep Learning. In *ICLR*.
- Spadaro, G.; Renzulli, R.; Bragagnolo, A.; Giraldo, J. H.; Fiandrotti, A.; Grangetto, M.; and Tartaglione, E. 2023. Shannon Strikes Again! Entropy-Based Pruning in Deep Neural Networks for Transfer Learning Under Extreme Memory and Computation Budgets. In *ICCVW*.
- Tartaglione, E.; Bragagnolo, A.; Fiandrotti, A.; and Grangetto, M. 2022. Loss-based sensitivity regularization: towards deep sparse neural networks. *Neural Networks*.
- Tishby, N.; Pereira, F.; Bialek, W.; Hajek, B.; and Sreenivas, R. 1999. Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing.
- Tishby, N.; and Zaslavsky, N. 2015. Deep learning and the information bottleneck principle. In *IEEE ITW*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *NeurIPS*.
- Wei, Z.; Hao, L.; and Zhang, X. 2022. Model Compression by Iterative Pruning with Knowledge Distillation and Its Application to Speech Enhancement. In *Interspeech*.
- Xu, K.; Rui, L.; Li, Y.; and Gu, L. 2020. Feature normalized knowledge distillation for image classification. In *ECCV*.
- Yilmaz, F. F.; and Heckel, R. 2022. Regularization-wise double descent: Why it occurs and how to eliminate it. In *ISIT*.
- Yim, J.; Joo, D.; Bae, J.; and Kim, J. 2017. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *CVPR*.
- Yuan, L.; Tay, F. E.; Li, G.; Wang, T.; and Feng, J. 2020. Revisiting knowledge distillation via label smoothing regularization. In *CVPR*.
- Zhang, Z.; Chen, X.; Chen, T.; and Wang, Z. 2021. Efficient lottery ticket finding: Less data is more. In *ICML*.
- Zhou, Z.; Zhou, Y.; Jiang, Z.; Men, A.; and Wang, H. 2022. An Efficient Method for Model Pruning Using Knowledge Distillation with Few Samples. In *ICASSP*.
- Zhu, M. H.; and Gupta, S. 2018. To Prune, or Not to Prune: Exploring the Efficacy of Pruning for Model Compression.