

AC/DC: Alternating Compressed/DeCompressed Training of Deep Neural Networks

Alexandra Peste* Eugenia Iofinova† Adrian Vladu‡ Dan Alistarh§

Abstract

The increasing computational requirements of deep neural networks (DNNs) have led to significant interest in obtaining DNN models that are *sparse*, yet *accurate*. Recent work has investigated the even harder case of *sparse training*, where the DNN weights are, for as much as possible, already sparse to reduce computational costs during training. Existing sparse training methods are mainly empirical and often have lower accuracy relative to the dense baseline. In this paper, we present a general approach called Alternating Compressed/DeCompressed (AC/DC) training of DNNs, demonstrate convergence for a variant of the algorithm, and show that AC/DC outperforms existing sparse training methods in accuracy at similar computational budgets; at high sparsity levels, AC/DC even outperforms existing methods that rely on accurate pre-trained dense models. An important property of AC/DC is that it allows *co-training* of dense and sparse models, yielding accurate *sparse–dense model pairs* at the end of the training process. This is useful in practice, where compressed variants may be desirable for deployment in resource-constrained settings without re-doing the entire training flow, and also provides us with insights into the accuracy gap between dense and compressed models.

1 Introduction

The tremendous progress made by deep neural networks in solving diverse tasks has driven significant research and industry interest in deploying efficient versions of these models. To this end, entire families of model compression methods have been developed, such as pruning [HABN⁺21] and quantization [GKD⁺21], which are now accompanied by hardware and software support [Van17, CMJ⁺18, DDJ⁺20, MLP⁺21, Gra21].

Neural network pruning, which is the focus of this paper, is the compression method with arguably the longest history [LDS90]. The basic goal of pruning is to obtain neural networks for which many connections are removed by being set to zero, while maintaining the network’s accuracy. A myriad pruning methods have been proposed—please see [HABN⁺21] for an in-depth survey—and it is currently understood that many popular networks can be compressed by more than an order of magnitude without significant accuracy loss.

Many accurate pruning methods require a fully-accurate, dense variant of the model, from which weights are subsequently removed. A shortcoming of this approach is the fact that the memory and computational savings due to compression are only available for the *inference*, post-training phase,

*IST Austria, alexandra.peste@ist.ac.at

†IST Austria, eugenia.iofinova@ist.ac.at

‡CNRS & IRIF, Université de Paris, vladu@irif.fr

§IST Austria, dan.alistarh@ist.ac.at

and not during training itself. This distinction becomes important especially for large-scale modern models, which can have millions or even billions of parameters, and for which fully-dense training can have high computational and even non-trivial environmental costs [SGM20].

One approach to address this issue is *sparse training*, which essentially aims to remove connections from the neural network as early as possible during training, while still matching, or at least approximating, the accuracy of the fully-dense model. For example, the RigL technique [EGM⁺20] randomly removes a large fraction of connections early in training, and then proceeds to optimize over the sparse support, providing savings due to sparse back-propagation. Periodically, the method re-introduces some of the weights during the training process, based on a combination of heuristics, which requires taking full gradients. These works, as well as many recent sparse training approaches [BKML17, MMS⁺18, JPR⁺20], which we cover in detail in the next section, have shown empirically that non-trivial computational savings, usually measured in theoretical FLOPs, can be obtained using sparse training, and that the optimization process can be fairly robust to sparsification of the support.

At the same time, this line of work still leaves intriguing open questions. The first is *theoretical*: to our knowledge, none of the methods optimizing over sparse support, and hence providing training speed-up, have been shown to have convergence guarantees. The second is *practical*, and concerns a deeper understanding of the relationship between the densely-trained model, and the sparsely-trained one. Specifically, (1) current sparse training methods can still leave a non-negligible accuracy gap, relative to dense training, or even post-training sparsification; and (2) most existing work on sparsity requires significant changes to the training flow, and focuses on maximizing global accuracy metrics; thus, we lack understanding when it comes to *co-training* sparse and dense models, as well as with respect to correlations between sparse and dense models *at the level of individual predictions*.

Contributions. In this paper, we take a step towards addressing these questions. We investigate a general approach for sparse training of neural networks, which we call *Alternating Compressed / DeCompressed (AC/DC)* training. AC/DC performs *co-training of sparse and dense models*, and can return both an accurate *sparse* model, and a *dense* model, which can recover the dense baseline accuracy via fine-tuning. We show that a variant of AC/DC ensures convergence for smooth, non-convex objectives, under analytic assumptions. Extensive experimental results show that it provides state-of-the-art accuracy among sparse training techniques at comparable training budgets, and can even outperform *post-training* sparsification approaches.

AC/DC builds on the classic *iterative hard thresholding (IHT)* family of methods for sparse recovery [BD08]. As the name suggests, AC/DC works by alternating the standard *dense* training phases, with *sparse* phases where optimization is performed exclusively over a fixed *sparse support*, and a subset of the weights and their gradients are fixed at zero, leading to computational savings. (This is in contrast to *error feedback* algorithms, e.g. [CHS⁺16, LSB⁺19] which require fully-dense gradients, even though the weights may be sparse.) The process uses the same number of epochs and hyper-parameters as regular training, and the frequency and length of the phases can be safely set to standard values, e.g. 5–10 epochs. We ensure that training ends on a *sparse* phase, and return the resulting *sparse* model, as well as the last *dense* model obtained at the end of a *dense* phase. This dense model may be additionally fine-tuned for a short period, leading to a more accurate *dense-finetuned* model, which we usually find to match the accuracy of the *dense baseline*.

We emphasize that algorithms alternating sparse and dense training phases for deep neural networks have been previously investigated [JYFY16, HPN⁺16], but with the different goal on using

sparsity as a regularizer to obtain *more accurate dense models*. Relative to these works, our goal is to produce highly-accurate, highly-sparse models, as well as maximizing the fraction of training time for which optimization is performed over a sparse support, leading to computational savings. Also, we are the first to provide convergence guarantees for variants of this approach.

We perform an extensive empirical investigation of this approach, showing that it provides consistently good results on a wide range of models and tasks (ResNet [HZRS16] and MobileNets [HZC⁺17] on the ImageNet [RDS⁺15] / CIFAR [KH⁺09] datasets, and Transformers [VSP⁺17, DYY⁺19] on WikiText [MXBS16]), under standard values of the training hyper-parameters. Specifically, when executed on the same number of training epochs, our method outperforms all previous *sparse training* methods in terms of the accuracy of the resulting sparse model, often by significant margins. This comes at the cost of slightly higher theoretical computational cost relative to prior sparse training methods, although AC/DC usually reduces training FLOPs to 45–65% of the dense baseline. AC/DC is also close to the accuracy of state-of-the-art post-training pruning methods [KRS⁺20, SA20] at medium sparsities (80% and 90%); surprisingly, it *significantly outperforms* them at higher sparsities. In addition, we showcase the flexibility of AC/DC by employing it to obtain *semi-structured* pruned models using the 2:4 sparsity pattern efficiently supported by new NVIDIA hardware [MLP⁺21]. Further, we show that the resulting sparse models can provide real-world speedups for DNN inference on CPUs [Dee21].

An interesting feature of AC/DC is that it allows for accurate dense/sparse co-training of models. Specifically, at medium sparsity levels (80% and 90%), the method allows the co-trained dense model to recover the dense baseline accuracy via a short fine-tuning period. In addition, dense/sparse co-training provides us with a lens into the training dynamics, in particular relative to the sample-level accuracy of the two models, but also in terms of the dynamics of the sparsity masks. Specifically, we observe that co-trained sparse/dense pairs have higher sample-level agreement than sparse/dense pairs obtained via post-training pruning, and that weight masks still change later in training.

Additionally, we probe the accuracy differences between sparse and dense models, by examining their “memorization” capacity [ZBH⁺16]. For this, we perform dense/sparse co-training in a setting where a small number of valid training samples have *corrupted labels*, and examine how these samples are classified during dense and sparse phases, respectively. We observe that the sparse model is less able to “memorize” the corrupted labels, and instead often classifies the corrupted samples to their *true* (correct) class. By contrast, during dense phases model can easily “memorize” the corrupted labels. (Please see Figure 3a for an illustration.) This suggests that one reason for the higher accuracy of dense models is their ability to “memorize” hard-to-classify samples.

2 Related Work

There has recently been tremendous research interest into pruning techniques for DNNs; we direct the reader to the recent surveys of [GEH19] and [HABN⁺21] for a comprehensive overview. Roughly, most DNN pruning methods can be split as (1) *post-training* pruning methods, which start from an accurate dense baseline, and remove weights, followed by fine-tuning; and (2) *sparse training* methods, which perform weight removal during the training process itself. (Other categories such as *data-free* pruning methods [LAT18, TKYG20] exist, but they are beyond our scope.) We focus on *sparse training*, although we will also compare against state-of-the-art post-training methods. Another perspective is *computational*: *post-training* methods produce efficient model *inference models*, but usually still compute dense gradients, while *sparse training* methods aim to also perform

sparse backpropagation, leading to computational gains.

Arguably, the most popular metric for weight removal is *weight magnitude* [Hag94, HPTD15, ZG17]. Better-performing approaches exist, such as second-order metrics [LDS90, HSW93, DCP17, SA20], or Bayesian approaches [MAV17], but these tend to have higher computational and implementation cost.

The general goal of *sparse training* methods is to perform both the forward (inference) pass *and the backpropagation* pass over a sparse support, leading to computational gains during the training process as well. One of the first approaches to maintain sparsity throughout training was Deep Rewiring [BKML17], where SGD steps applied to positive weights are augmented with random walks in parameter space, followed by inactivating negative weights. To maintain sparsity throughout training, randomly chosen inactive connections are re-introduced in the “growth” phase. Sparse Evolutionary Training (SET) [MMS⁺18] introduces a non-uniform sparsity distribution across layers, which scales with the number of input and output channels, and trains sparse networks by pruning weights with smallest magnitude and re-introducing some weights randomly. RigL [EGM⁺20], prunes weights at random after a warm-up period, and then periodically performs weight re-introduction using a combination of connectivity- and gradient-based statistics, which require periodically evaluating full gradients. RigL can lead to state-of-the-art accuracy results even compared to post-training methods; however, for this it requires significant additional data passes (e.g. 5x) relative to the dense baseline. Most recently, Top-KAST [JPR⁺20] alleviated the drawback of periodically having to evaluate dense gradients, by updating the sparsity masks using gradients *of reduced sparsity* relative to the weight sparsity. This allows some of the pruned weights to be re-introduced during training, improving accuracy. The latter two methods set the state-of-the-art for sparse training: when executing for the same number of epochs as the dense baseline, they provide computational reductions the order of 2x, while the accuracy of the resulting sparse models is lower than that of leading post-training methods, executed at the same sparsity levels. To our knowledge, all these methods are heuristics, i.e. do not have convergence guarantees.

Another approach towards pruning with an intermediate computational cost is *training sparse networks from scratch*. The masks are updated by continuously pruning and re-introducing weights. For example, [LSB⁺19] uses magnitude pruning after applying SGD on the dense network, whereas [DZ19] update the masks by re-introducing weights with the highest gradient momentum. STR [KRS⁺20] learns a separate pruning threshold for each layer and allows sparsity both during forward and backward passes; however, the desired sparsity can not be explicitly imposed, and the network has low sparsity for a large portion of training. These methods can lead to only limited computational gains, since they either require dense gradients, or the sparsity level cannot be imposed. By comparison, our method provides models of similar or better accuracy at the same sparsity, with computational reductions. We also obtain dense models that match the baseline accuracy, with a fraction of the baseline FLOPS.

The idea of alternating sparse and dense training phases has been examined before in the context of neural networks, but with the goal of using temporary sparsification as a regularizer. Specifically, *Dense-Sparse-Dense (DSD)* [HPN⁺16] proposes to first *train a dense model to full accuracy*; this model is then sparsified via magnitude; next, optimization is performed over the sparse support, followed by an additional optimization phase over the full dense support. Thus, this process is used as a regularization mechanism for the dense model, which results in relatively small, but consistent accuracy improvements relative to the original dense model. In [JYFY16], the authors

propose a similar approach to DSD, but alternate sparse phases during the regular training process. The resulting process is similar to AC/DC, but, importantly, the goal of their procedure is to return a *more accurate dense model*. (Please see their Algorithm 1.) For this, the authors use relatively low sparsity levels, and gradually increase sparsity during optimization; they observe accuracy improvements for the resulting dense models, at the cost of increasing the total number of epochs of training. By contrast, our focus is on obtaining accurate *sparse* models, while reducing computational cost, and executing the dense training recipe. We execute at higher sparsity levels, and on larger-scale datasets and models. In addition, we also show that the method works for other sparsity patterns, e.g. the 2:4 semi-structured pattern [MLP⁺21].

More broadly, the Lottery Ticket Hypothesis (LTH) [FC18], states that sparse networks can be trained in isolation *from scratch* to the same performance as a post-training pruning baseline, by starting from the “right” weight and sparsity mask initializations, optimizing only over this sparse support. However, initializations usually require the availability of the fully-trained dense model, falling under *post-training* methods. There is still active research on replicating these intriguing findings to large-scale models and datasets [GEH19, FDRC19].

3 Alternating Compressed / DeCompressed (AC/DC) Training

3.1 Background and Assumptions

Obtaining *sparse* solutions to optimization problems is a problem of interest in several areas [CT06, BD08, Fou11], where the goal is to minimize a function $f : \mathbb{R}^N \rightarrow \mathbb{R}$ under sparsity constraints:

$$\min_{\theta \in \mathbb{R}^N} f(\theta) \quad \text{s.t.} \quad \|\theta\|_0 \leq k. \quad (1)$$

For the case of ℓ_2 regression, $f(\theta) = \|b - A\theta\|_2^2$, a solution has been provided by Blumensath and Davies [BD08], known as the *Iterative Hard Thresholding (IHT)* algorithm, and subsequent work [Fou11, Fou12, YLZ14] provided theoretical guarantees for the linear operators used in compressed sensing. The idea consists of alternating gradient descent (GD) steps and applications of a thresholding operator to ensure the ℓ_0 constraint is satisfied. More precisely, T_k is defined as the “top- k ” operator, which keeps the largest k entries of a vector θ in absolute value, and replaces the rest with 0. The IHT update at step $t + 1$ has the following form:

$$\theta_{t+1} = T_k(\theta_t - \eta \nabla f(\theta_t)). \quad (2)$$

Most convergence results for IHT assume deterministic gradient descent steps. For DNNs, stochastic methods are preferred, so we describe and analyze a stochastic version of IHT and prove its convergence.

Stochastic IHT. We consider functions $f : \mathbb{R}^N \rightarrow \mathbb{R}$, for which we can compute stochastic gradients g_θ , which are unbiased estimators of the true gradient $\nabla f(\theta)$. Define the *stochastic* IHT update as:

$$\theta_{t+1} = T_k(\theta_t - \eta g_{\theta_t}). \quad (3)$$

This formulation covers the practical case where the stochastic gradient g_θ corresponds to a *mini-batch* stochastic gradient. Indeed, as in practice f takes the form $f(\theta) = \frac{1}{m} \sum_{i=1}^m f(\theta; x_i)$, where $S = \{x_1, \dots, x_m\}$ are data samples, the stochastic gradients obtained via backpropagation take the form $\frac{1}{|B|} \sum_{i \in B} \nabla f(\theta; x_i)$, where B is a sampled mini-batch. We aim to prove strong convergence bounds for stochastic IHT, under common assumptions that arise in the context of training DNNs.

Analytical Assumptions. Formally, our analysis uses the following assumptions on f .

1. Unbiased gradients with variance σ : $\mathbb{E}[g_\theta|\theta] = \nabla f(\theta)$, and $\mathbb{E}[\|g_\theta - \nabla f(\theta)\|^2] \leq \sigma^2$.
2. Existence of a k^* -sparse minimizer θ^* : $\exists \theta^* \in \arg \min_\theta f(\theta)$, s.t. $\|\theta^*\|_0 \leq k^*$.
3. For $\beta > 0$, the β -smoothness condition when restricted to t coordinates ((t, β) -smoothness):

$$f(\theta + \delta) \leq f(\theta) + \nabla f(\theta)^\top \delta + \frac{\beta}{2} \|\delta\|^2, \text{ for all } \theta, \delta \text{ s.t. } \|\delta\|_0 \leq t. \quad (4)$$

4. For $\alpha > 0$ and number of indices r , the r -concentrated Polyak-Łojasiewicz $((r, \alpha)$ -CPL) condition:

$$\|T_r(\nabla f(\theta))\| \geq \frac{\alpha}{2} (f(\theta) - f(\theta^*)), \text{ for all } \theta. \quad (5)$$

The first assumption is standard in stochastic optimization, while the existence of very sparse minimizers is a known property in over-parametrized DNNs [FC18], and is the very premise of our study. Smoothness is also a standard assumption, e.g. [LSB⁺19]—we only require it along *sparse* directions, which is a strictly weaker assumption. The more interesting requirement for our convergence proof is the (r, α) -CPL condition in Equation (5).

The standard Polyak-Łojasiewicz (PL) condition [KNS16] is common in non-convex optimization, and versions of it are essential in the analysis of DNN training [LZB20, AZLS19]. Its standard form states that small gradient norm, i.e. approximate stationarity, implies closeness to optimum in function value. We require a slightly stronger version, in terms of the norm of the gradient contributed by its largest coordinates in absolute value. This restriction appears necessary for the success of IHT methods, as the sparsity enforced by the truncation step automatically reduces the progress ensured by a gradient step to an amount proportional to the norm of the top- k gradient entries. This strengthening of the PL condition is supported both theoretically, by the mean-field view, which argues that gradients are sub-gaussian [SM20], and by empirical validations of this behaviour [AHJ⁺18, SCCS19].

We are now ready to state our main analytical result.

Theorem 1. *Let $f : \mathbb{R}^N \rightarrow \mathbb{R}$ be a function with a k^* -sparse minimizer θ^* . Let $\beta > \alpha > 0$ be parameters, let $k = C \cdot k^* \cdot (\beta/\alpha)^2$ for some appropriately chosen constant C , and suppose that f is $(2k + 3k^*, \beta)$ -smooth and (k^*, α) -CPL. For initial parameters θ_0 and precision $\epsilon > 0$, given access to stochastic gradients with variance σ , stochastic IHT (3) converges in $O\left(\frac{\beta}{\alpha} \cdot \ln \frac{f(\theta_0) - f(\theta^*)}{\epsilon}\right)$ iterations to a point θ with $\|\theta\|_0 \leq k$, such that*

$$\mathbb{E}[f(\theta) - f(\theta^*)] \leq \epsilon + \frac{16\sigma^2}{\alpha}.$$

Roughly, this result suggests that, for a given set of parameters, by increasing the support demanded from our approximate minimizer θ relative to the optimal k^* , the parameter α increases, which leads to lower loss, and potentially to faster running time. We provide a complete proof of this result in Appendix A. Our analysis approach also works in the absence of the CPL condition (Theorem 3), in which case we prove that a version of the algorithm can find sparse nearly-stationary points. As a bonus, we also simplify existing analyses for IHT and extend them to the stochastic case (Theorem 2). Another interpretation of our results is in showing that, under our assumptions, *error feedback* [LSB⁺19] is not necessary for recovering good sparse minimizers; this has practical implications, as it allows us to perform fully-sparse back-propagation in sparse optimization phases. Next, we discuss our practical implementation, and its connection to these theoretical results.




Figure 1: The AC/DC training process. After a short warmup we alternatively prune to maximum sparsity and restore the pruned weights. The plot shows the sparsity and validation accuracy throughout the process for a sample run on ResNet50/ImageNet at 90% sparsity.

Algorithm 1 Alternating Compressed/Decompressed (AC/DC) Training

Require: Weights $\theta \in \mathbb{R}^N$, data S , sparsity k , compression phases \mathcal{C} , decompression phases \mathcal{D}

- 1: Train the weights θ for Δ_w epochs ▷ Warm-up phase
- 2: **while** epoch \leq max epochs **do**
- 3: **if** entered a compression phase **then**
- 4: $\theta \leftarrow T_k(\theta, k)$ ▷ apply compression (top-k) operator on weights
- 5: $m \leftarrow \mathbb{1}[\theta_i \neq 0]$ ▷ create masks
- 6: **end if**
- 7: **if** entered a decompression phase **then**
- 8: $m \leftarrow \mathbb{1}_N$ ▷ reset all masks
- 9: **end if**
- 10: $\theta \leftarrow \theta \odot m$ ▷ apply the masks (ensure sparsity for compression phases)
- 11: $\tilde{\theta} \leftarrow \{\theta_i | m_i \neq 0, 1 \leq i \leq N\}$ ▷ get the support for the gradients
- 12: **for** x mini-batch in S **do**
- 13: $\theta \leftarrow \theta - \eta \nabla_{\tilde{\theta}} f(\theta; x)$ ▷ optimize the active weights
- 14: **end for**
- 15: epoch \leftarrow epoch + 1
- 16: **end while**
- 17: **return** θ

3.2 AC/DC: Applying IHT to Deep Neural Networks

AC/DC starts from a standard DNN training flow, using standard optimizers such as SGD with momentum [Qia99] or Adam [KB14] for a fixed number of epochs. It preserves the standard training hyper-parameters, e.g., learning rate, momentum, weight decay, and will only periodically modify the *support* for optimization. Please see Algorithm 1 for pseudocode.

We partition the set of training epochs into *compressed* epochs \mathcal{C} , and *decompressed* epochs \mathcal{D} . We begin with a *dense warm-up* period of Δ_w consecutive epochs, during which regular dense (decompressed) training is performed. We then start alternating *compressed optimization* phases of length Δ_c epochs each, with *decompressed (regular) optimization* phases of length Δ_d epochs each. The process completes on a compressed fine-tuning phase, returning an accurate sparse model. Alternatively, if our goal is to return a dense model matching the baseline accuracy, we take the best dense checkpoint obtained during alternation, and fine-tune it over the entire support. In practice, we noticed that allowing a longer final decompressed phase of length $\Delta_D > \Delta_d$ improves the performance of the dense model, by allowing it to better recover the baseline accuracy after fine-tuning. Please see Figure 1 for an illustration of the schedule.

In our experiments, we focus on the case where the compression operation is unstructured or semi-structured pruning. In this case, at the beginning of each sparse optimization phase, we apply

the top-k operator across all of the network weights to obtain a mask M over the weights θ . The top-k operator is applied globally across all of the network weights, and will represent the sparse support over which optimization will be performed for the rest of the current sparse phase. At the end of the sparse phase, the mask M is reset to all-1s, so that the subsequent dense phase will optimize over the full dense support. Furthermore, once all weights are re-introduced, it is beneficial to reset to 0 the gradient momentum term of the optimizer; this is particularly useful for the weights that were previously pruned, which would otherwise have stale versions of gradients.

Discussion. Moving from the theoretical algorithm to a robust practical implementation required some adjustments, so we now interpret AC/DC from the perspective of our analysis. First, each *decompressed phase* can be directly mapped to a *deterministic/stochastic IHT* step, where, instead of a single gradient step in between consecutive truncations of the support, we perform several stochastic steps. These additional steps improve the accuracy of the method in practice, and we can bound their influence in theory as well, although they do not necessarily provide better bounds. This leaves open the interpretation of the *compressed phases*: for this, notice that the core of the proof for Theorem 1 is in showing that a single IHT step significantly decreases the expected value of the objective; using a similar argument, we can prove that additional optimization steps over the sparse support can only improve convergence. We are able to show convergence for a variant of IHT closely following AC/DC, but the bounds do not improve over Theorem 1.

4 Experimental Validation

Goals and Setup. We tested AC/DC on image classification tasks (CIFAR-100 [KH⁺09] and ImageNet [RDS⁺15]) and on language modelling tasks [MXBS16] using the Transformer-XL model [DYY⁺19]. The goal is to examine the *validation accuracy* of the resulting sparse and dense models, versus the induced sparsity, as well as the number of FLOPs used for training and inference, relative to other sparse training methods. Additionally, we compare against state-of-the-art post-training pruning methods [SA20, KRS⁺20]. We also examine prediction differences between the sparse and dense models.

We use PyTorch [PGM⁺19] for our implementation, and Weights & Biases [Bie20] for experimental tracking. Experiments were run on NVIDIA RTX 2080 GPUs for image classification tasks, and NVIDIA RTX 3090 GPUs for language modelling. Each ImageNet run took approximately 2 days for ResNet50 and one day for MobileNet, while each Transformer-XL experiment took approximately 2 days. All reported image classification experiments were performed in triplicate by varying the random seed; we report mean and standard deviation. Due to computational limitations, the language modelling experiments were conducted in a single run.

4.1 CIFAR-100 Experiments

We tested the IHT pruning algorithm on the CIFAR-100 dataset [KH⁺09]. We used the WideResNet-28-10 architecture [ZK16], which to our knowledge gives state-of-the-art performance on this dataset. Models were trained for a fixed period of 200 epochs, using Stochastic Gradient Descent (SGD) with momentum, and a stepwise decreasing learning rate scheduler. After a warm-up period of ten epochs, we alternated sparse and dense phases of 20 epochs until epoch 170, at which point we allowed the sparse model to train to convergence for 30 more epochs. In addition, we found a small benefit to resetting momentum to 0 at each transition from sparse to dense, and we have done so throughout the trials. All experiments were replicated starting from three different seeds, and we report average results and their standard deviations.

Table 1: CIFAR-100 Sparsity results on WideResNet

Method	Top-1 Acc. (%)	Sparsity	GFLOPs Inference	EFLOPs Train
Dense	79.0 ± 0.25	0%	11.9	0.36
AC/DC	79.6 ± 0.17	49.98%	0.50×	0.72×
GMP	79.2 ± 0.17	49.98%	0.46×	1.64×
AC/DC	80.0 ± 0.17	74.96%	0.29×	0.6×
GMP	78.9 ± 0.14	74.96%	0.26×	1.52×
AC/DC	79.1 ± 0.07	89.96%	0.114×	0.51×
GMP	77.7 ± 0.23	89.96%	0.13×	1.44×
AC/DC	78.2 ± 0.12	94.95%	0.08×	0.47×
GMP	76.6 ± 0.07	94.95%	0.07×	1.41×

We compare our results with Gradual Magnitude Pruning [ZG17]. To our knowledge, we are the first to release CIFAR-100 pruning results for this network architecture, and GMP was chosen as a baseline due to its generally strong performance against a range of other approaches [GEH19]. We obtain the GMP baseline by training on the WideResNet-28-10 architecture at full density for 50 epochs, then gradually increasing the sparsity over the next 100 before allowing the model to converge for the final 50, matching the 200 training epochs of AC/DC. We further validated this baseline by training dense WideResnet-28-10 models for 200 epochs and then gradually pruning over 50 and finetuning over 50 more, for a total of 300 epochs, which gave similar performance at the cost of greater FLOPs and training time.


The results are shown in Table 1. We see that AC/DC pruning significantly outperforms Gradual Magnitude Pruning at all sparsity levels tested, and further that AC/DC models pruned at 50% and 75% even outperform the dense baseline, while the models pruned at 90% at least match it.

4.2 ImageNet Experiments


On the ImageNet dataset [RDS⁺15], we test AC/DC on ResNet50 [HZRS16] and MobileNetV1 [HZC⁺17]. In all reported results, the models were trained for a fixed number of 100 epochs, using SGD with momentum. We use a cosine learning rate scheduler and training hyper-parameters following [KRS⁺20], but without label smoothing. Our dense baseline has 76.84% validation accuracy. We report the average and the standard deviation across 3 different seeds. Unless otherwise specified, weights are pruned globally, based on their magnitude and in a single step. Similar to previous work, we did not prune biases, nor the Batch Normalization parameters.

For all results, the AC/DC compression schedule starts with a “warm-up” phase of dense training for 10 epochs, after which we alternate between compression and de-compression every 5 epochs, until the last dense and sparse phase. It is beneficial to allow these last two “fine-tuning” phases to run longer: the last decompression phase runs for 10 epochs, whereas the final 15 epochs are the compression fine-tuning phase. We reset SGD momentum at the beginning of every decompression phase. In total, we have an equal number of epochs of dense and sparse training; see Figure 2a for an illustration. We use exactly the same setup for both ResNet50 and MobileNetV1 models, which resulted in high-quality sparse models. To recover a dense model with baseline accuracy using AC/DC, we finetune the best dense checkpoint obtained during training; practically, this replaces the last *sparse* fine-tuning phase with a phase where the *dense* model is fine-tuned instead.

ResNet50 Results. Tables 2 and 3 contain validation accuracy results across medium and high global sparsity levels, as well as inference and training FLOPs. Overall, AC/DC achieves higher



(a) Sparsity pattern and test accuracy



(b) Relative change in consecutive masks

Figure 2: Validation accuracy and sparsity during training, together with differences in consecutive masks for ResNet50 on ImageNet using AC/DC.

validation accuracy than any of the state-of-the-art sparse training methods, when using the same number of epochs. At the same time, due to dense training phases, AC/DC tends to have higher FLOP requirements relative to RigL or Top-KAST at the same sparsity.

Table 2: Results for medium sparsity on ResNet50.

Method	Sparsity (%)	Top-1 Acc. (%)	GFLOPs Inference	EFLOPs Train
Dense	0	76.84	8.2	3.15
AC/DC	79.83	76.3 \pm 0.1	0.29 \times	0.65 \times
RigL ₁ \times (ERK)	80	75.1 \pm 0.05	0.42 \times	0.42 \times
Top-KAST	80	75.03	0.23 \times	0.32 \times
STR	79.55	76.19	0.19 \times	-
WoodFisher	80	76.76	0.25 \times	-
AC/DC	89.8	75.03 \pm 0.1	0.19 \times	0.58 \times
RigL ₁ \times (ERK)	90	73.0 \pm 0.04	0.25 \times	0.24 \times
Top-KAST	90	74.76	0.13 \times	-
STR	90.23	74.31	0.08 \times	-
WoodFisher	90	75.21	0.14 \times	-

Table 3: Results for high sparsity on ResNet50.

Method	Sparsity (%)	Top-1 Acc. (%)	GFLOPs Inference	EFLOPs Train
Dense	0	76.84	8.2	3.15
AC/DC	94.8	73.14 \pm 0.2	0.12 \times	0.5 \times
RigL ₁ \times (ERK)	95	69.7 \pm 0.17	0.12 \times	0.42 \times
Top-KAST	95	70.42	0.08 \times	0.1 \times
STR	94.8	70.97	0.04 \times	-
WoodFisher	95	72.12	-	-
AC/DC	97.8	68.44 \pm 0.09	0.06 \times	0.46 \times
Top-KAST	98	67.06	0.05 \times	0.08 \times
STR	97.78	62.84	0.02 \times	-
WoodFisher	98	65.55	-	-

At medium sparsities (80% and 90%), AC/DC sparse models are slightly less accurate than the state-of-the-art post-training methods (WoodFisher), by small margins. The situation is reversed at higher sparsities, where AC/DC produces significantly more accurate models: the gap to the second-best methods (WoodFisher / Top-KAST) is of more than 1% at 95% and 98% sparsity.

Of the existing sparse training methods, Top-KAST is closest in terms of validation accuracy to our sparse model, at 90% sparsity. However, Top-KAST does not prune the first and last layers, whereas the results in the tables do not restrict the sparsity pattern. For fairness, we executed AC/DC using the same layer-wise sparsity distribution as Top-KAST, for both uniform and global magnitude pruning. For 90% global pruning, results for AC/DC improved significantly; the best sparse model reached 75.63% validation accuracy (0.6% increase over Table 2), while the best dense model had 76.85% after fine-tuning. For uniform sparsity, our results did not change significantly: 75.04% validation accuracy for the sparse model and 76.43% - for the fine-tuned dense model. We also note that Top-KAST has better results at 98% when increasing the number of training epochs 2 times, and considerably fewer training FLOPs (e.g. 15% of the dense FLOPs). For fairness, we compared against all methods on a fixed number of 100 training epochs and we additionally trained AC/DC at high sparsity without pruning the first and last layers. Our results improved significantly: 74.16% accuracy for 95% sparsity, and 71.27% for 98% sparsity, both surpassing Top-KAST with prolonged training. We provide a more detailed comparison in Appendix B.

An advantage of AC/DC is that it provides *both* sparse and dense models at cost *below* that of a single dense training run. For medium sparsity, the accuracy of the dense-finetuned model is very close to the dense baseline. Concretely, at 90% sparsity, with 58% of the total (theoretical) baseline training FLOPs, we obtain a *sparse* model which is close to state of the art; in addition, by fine-tuning the best dense model, we obtain a dense model with 76.56% (average) validation accuracy. The whole process takes 73% of the baseline training FLOPs.

Interestingly, the sparsity distribution over layers does not change significantly during training; yet, the dynamics of the masks has an important impact on the performance of AC/DC, as it can enable faster recovery from sub-optimal pruning decisions. Specifically, we observed that masks keep getting updated over time, although the change between consecutive sparse masks decreases (please see Figure 2b), and small percentage of the weights remain fixed at 0 even during dense training. Please see Appendix B for additional results and analysis.

Table 4: Sparsity results on MobileNetV1.

Method	Sparsity (%)	Top-1 Acc. (%)	GFLOPs Inference	EFLOPs Train
Dense	0	71.78	1.1	0.44
AC/DC	74.61	70.3 \pm 0.07	0.35 \times	0.65 \times
RigL $_{1\times}$ (ERK)	75	68.39	0.52 \times	0.53 \times
STR	75.28	68.35	0.18 \times	-
WoodFisher	75.28	70.09	0.28 \times	-
AC/DC	89.52	66.08 \pm 0.09	0.19 \times	0.57 \times
RigL $_{1\times}$ (ERK)	90	63.58	0.27 \times	0.29 \times
STR	89.01	62.1	0.07 \times	-
WoodFisher	89	63.87	0.15 \times	-

Table 5: Sparsity results on Transformer-XL.

Method	Sparsity (%)	Perplexity Sparse	Perplexity Dense	Perplexity Finetuned Dense
Dense	0	-	18.95	-
AC/DC	80	20.93	20.36	19.46
AC/DC	80, 50 embed.	20.99	20.28	19.5
Top-KAST	80, 0 bwd	19.8	-	-
Top-KAST	80, 60 bwd	21.3	-	-
AC/DC	90	22.65	21.07	20.08
AC/DC	90, 50 embed.	22.81	21.02	20.07
Top-KAST	90, 80 bwd	25.1	-	-

MobileNet Results. We perform the same experiment, using exactly the same setup, on the MobileNetV1 architecture [HZC⁺17], which is compact and thus harder to compress. On a training budget of 100 epochs, our method finds sparse models with higher Top-1 validation accuracy than existing sparse- and post-training methods, on both 75% and 90% sparsity levels (Table 4). Importantly, AC/DC uses exactly the same hyper-parameters used for training the dense baseline [KRS⁺20]. Similar to ResNet50, at 75% sparsity, the *dense-finetuned model* recovers the baseline performance, while for 90% it is less than 1% below the baseline.

The only method which obtains higher accuracy for the same sparsity is the version of RigL [EGM⁺20] which executes for 5x more training epochs than the dense baseline. However, this version also uses more computation than the dense model. We limit ourselves to a fixed number of 100 epochs, the same used to train the dense baseline, which would allow for savings in training time. Moreover, RigL does not prune the first layer and the depth-wise convolutions, whereas for the results reported we do not impose any sparsity restrictions. Overall, we found that keeping these layers dense improved our results on 90% sparsity by almost 0.5%. Then, our results are quite close to RigL $_{2\times}$, with half the training epochs, and less training FLOPs. We provide a more detailed comparison in Appendix B.2.

Semi-structured Sparsity. We also experiment with the recent 2:4 sparsity pattern (2 weights out of each block of 4 are zero) proposed by NVIDIA, which ensures inference speedups on the Ampere architecture. Recently, [MLP⁺21] showed that accuracy can be preserved under this pattern, by deriving masks from the pretrained dense model, and re-doing the entire training flow. Also, [ZMZ⁺21] proposed more general N:M structures, together with a method for training such sparse models from scratch. We applied AC/DC to the 2:4 pattern, performing training from scratch and




Figure 3: Accuracy during training with AC/DC at various target sparsities, for 1000 randomly labelled CIFAR10 images, w.r.t their true labels.

obtained sparse models with $76.64\% \pm 0.05$ validation accuracy, i.e. slightly below the baseline. Furthermore, the dense-finetuned model fully recovers the baseline performance (76.85% accuracy).

4.3 Language Modeling

Next, we apply AC/DC to compressing NLP models. We use Transformer-XL [DYY⁺19], on the WikiText-103 dataset [MXBS16], with the standard model configuration with 18 layers and 285M parameters, trained using the Lamb optimizer [YLR⁺19] and standard hyper-parameters, which we describe in Appendix C.2.

The same Transformer-XL model trained on WikiText-103 was used in Top-KAST [JPR⁺20], which allows a direct comparison. Similar to Top-KAST, we did not prune the embedding layers, as this greatly affects the quality, without reducing computational cost. (For completeness, we do provide results when embeddings are pruned to 50% sparsity.) Our sparse training configuration consists in starting with a dense warm-up phase of 5 epochs, followed by alternating between compression and decompression phases every 3 epochs; we follow with a longer decompression phase between epochs 33-39, and end with a compression phase between epochs 40-48.

The results are shown in Table 5. Relative to Top-KAST, our approach provides significantly improved test perplexity at 90% sparsity, as well as better results at 80% sparsity with sparse back-propagation. (Top-KAST provides superior perplexity at 0% backward sparsity, but this provides little to no training speedup.) The results confirm that AC/DC is scalable and extensible. We note that our hyper-parameter tuning for this experiment was minimal.

4.4 Additional Properties

Output Analysis. We probe the accuracy difference between the sparse and dense-finetuned models. We first examined *sample-level agreement* between sparse and dense-finetuned pairs produced by AC/DC, relative to model pairs produced by gradual magnitude pruning (GMP). Co-trained model pairs consistently agree on more samples relative to GMP: for example, on the 80%-pruned ResNet50 model, the AC/DC model pair agrees on the Top-1 classification of 90% of validation samples, whereas the GMP models agree on 86% of the samples. A similar trend holds for the *cross-entropy* between model outputs.

Memorization. Second, we analyze differences in “memorization” capacity [ZBH⁺16] between dense and sparse models. For this, we apply AC/DC to ResNet20 trained on a variant of CIFAR-10 where a subset of 1000 samples have randomly corrupted class labels, and examine the accuracy on these samples during training. We first consider the setting where no data augmentation is used, and study

90% and 95% sparsity AC/DC runs. Figure 3a shows the results, when the accuracy for each sample is measured with respect to the *true, un-corrupted* label. During early training and during *sparse phases*, the network tends to classify corrupted samples to their *true class*, “ignoring” label corruption. However, as training progresses, due to dense training phases and lower learning rate, networks tend to “memorize” these samples, assigning them to their corrupted class. This phenomenon is even more prevalent at 95% sparsity, where the network is less capable of memorization. Interestingly, when data augmentation is used during training, both the sparse and dense models correctly classify most of the perturbed samples to their true labels, in the case of AC/DC runs at higher sparsity levels (90% and 95%). For lower sparsity levels (50% and 75%), more of the perturbed samples are memorized. This behavior can be observed in Figure 3b. We discuss these findings in more detail in Appendix B.5.

Practical Speedups. One remaining question regards the potential of sparsity to provide real-world speedups. While this is an active research area, e.g. [EDGS20], we partially address this concern in Appendix B.3, by showing inference speedups for our models on a CPU inference platform supporting unstructured sparsity [Dee21]: for example, our 90% sparse ResNet50 model provides 1.75x speedup for real-time inference (batch-size 1) on a resource-constrained processor with 4 cores, and 2.75x speedup on 16 cores at batch size 64, versus the dense model.

5 Conclusion, Limitations, and Future Work

We introduced AC/DC—a method for co-training sparse and dense models, with theoretical guarantees. Experimental results show that AC/DC improves upon the accuracy of previous sparse training methods, and obtains state-of-the-art results at high sparsities. Importantly, we recover near-baseline performance for dense models and do not require extensive hyper-parameter tuning. We also show that AC/DC has potential for real-world speed-ups in inference and training, with the appropriate software and hardware support. The method has the advantage of returning both an accurate standard model, and a compressed one, which may be useful in deployment scenarios. Our model output analysis confirms the intuition that sparse training phases act as a regularizer, preventing the (dense) model from memorizing corrupted samples. At the same time, they prevent the memorization of *hard samples*, which can affect accuracy.

The main limitations of AC/DC are its reliance on dense training phases, which limits the achievable training speedup, and the need for tuning the length and frequency of sparse/dense phases. We believe the latter issue can be addressed with more experimentation, as simple patterns tend to work well out-of-the-box; however, both the theoretical results and the output analysis suggest that dense phases may be *necessary* for good accuracy. We plan to further investigate this in future work, together with applying AC/DC to other compression methods, such as quantization, as well as leveraging sparse training on hardware that could efficiently support it, such as Graphcore IPUs [Gra21].

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 805223 ScaleML), and a CNRS PEPS grant. We would also like to thank Christoph Lampert for his feedback on an earlier version of this work, as well as for providing hardware for the Transformer-XL experiments.

Appendices

Table of Contents

1	Introduction	1
2	Related Work	3
3	Alternating Compressed / DeCompressed (AC/DC) Training	5
3.1	Background and Assumptions	5
3.2	AC/DC: Applying IHT to Deep Neural Networks	7
4	Experimental Validation	8
4.1	CIFAR-100 Experiments	8
4.2	ImageNet Experiments	9
4.3	Language Modeling	12
4.4	Additional Properties	12
5	Conclusion, Limitations, and Future Work	13
	Appendices	14
A	Convergence Proofs	15
A.1	Overview	15
A.2	Stochastic IHT for Non-Convex Functions with Concentrated PL Condition	18
A.3	Stochastic IHT for Functions with Restricted Smoothness and Strong Convexity	22
A.4	Finding a Sparse Nearly-Stationary Point	24
A.5	Deferred Proofs	25
B	Additional Experiments	29
B.1	ResNet50 on ImageNet	29
B.2	MobileNet on ImageNet	32
B.3	Inference Speedups	34
B.4	Sparse-Dense Output Comparison	35
B.5	Memorization Experiments on CIFAR10	36
C	Computational Details	37
C.1	FLOPs Computation	37

A Convergence Proofs

In this section we provide the convergence analysis for our algorithms. We prove Theorem 1 and show as a corollary that under reasonable assumptions our implementation of AC/DC converges to a sparse minimizer.

A.1 Overview

We use the notation and assumptions defined in Section 3.1. As all of our analyses revolve around bounding the progress made in a single iteration, to simplify notation we will generally use θ to denote the current iterate, and θ' to denote the iterate obtained after the IHT update:

$$\theta' = T_k(\theta - \eta g_\theta).$$

Additionally, we let $S, S', S^* \subseteq [N]$ to denote the support of θ, θ' and θ^* respectively, where θ^* is the promised k^* -sparse minimizer. Given an arbitrary vector x , we use $\text{supp}(x)$ to denote its support, i.e. the set of coordinates where x is nonzero. We may also refer to the minimizing value of f as $f^* = f(\theta^*)$.

Before providing more intuition for the analysis, we give the main theorem statements.

Stochastic IHT for Functions with Concentrated PL Condition. We restate the Theorem from Section 3.1.

Theorem 1. *Let $f : \mathbb{R}^N \rightarrow \mathbb{R}$ be a function with a k^* -sparse minimizer θ^* . Let $\beta > \alpha > 0$ be parameters, let $k = C \cdot k^* \cdot (\beta/\alpha)^2$ for some appropriately chosen constant C , and suppose that f is $(2k + 3k^*, \beta)$ -smooth and (k^*, α) -CPL. For initial parameters θ_0 and precision $\epsilon > 0$, given access to stochastic gradients with variance σ , stochastic IHT (3) converges in $O\left(\frac{\beta}{\alpha} \cdot \ln \frac{f(\theta_0) - f(\theta^*)}{\epsilon}\right)$ iterations to a point θ with $\|\theta\|_0 \leq k$, such that*

$$\mathbb{E}[f(\theta) - f(\theta^*)] \leq \epsilon + \frac{16\sigma^2}{\alpha}.$$

Additionally, we give a corollary that justifies our implementation of AC/DC. As opposed to the theoretical stochastic IHT algorithm, AC/DC performs a sequence of several dense SGD steps before applying a single pruning step. We show that even with this change we can provide theoretical convergence bounds, although these bounds can be weaker than the baseline IHT method under our assumptions.

Corollary 1 (Convergence of AC/DC). *Let $f : \mathbb{R}^N \rightarrow \mathbb{R}$ be a function that decomposes as $f(\theta) = \frac{1}{m} \sum_{i=1}^m f_i(\theta)$, and has a k^* -sparse minimizer θ^* . Let $\beta > \alpha > 0$ be parameters, let $k = C \cdot k^* \cdot (\beta/\alpha)^2$ for some appropriately chosen constant C , suppose that each f_i is (N, β) -smooth, and L -Lipschitz, and that f is (k^*, α) -CPL.*

Let Δ_c and B be integers, and let $\{S_1, \dots, S_B\}$ be a partition of $[m]$ into B subsets of cardinality $O(m/B)$ each. Given θ , let $g_\theta^{(i)} = \frac{1}{|S_i|} \sum_{j \in S_i} \nabla f_j(\theta)$.

Suppose we replace the IHT iteration with Δ_c phases during each of which we perform a full pass over the data by performing the iteration $\theta' = \theta - \eta g_\theta^{(i)}$ for all $i \in [B]$, and at the end of the Δ_c phases we apply the truncation operator T_k . This is optionally followed by a sparse training phase which fully optimizes f over the sparse support.

For initial parameters θ_0 and precision $\epsilon > 0$, this algorithm converges in $O\left(\frac{\beta}{\alpha} \cdot \ln \frac{f(\theta_0) - f(\theta^*)}{\epsilon}\right)$ iterations to a point θ with $\|\theta\|_0 \leq k$, such that

$$f(\theta) - f(\theta^*) \leq \epsilon + O\left(\frac{L^2}{\alpha}\right).$$

We provide the main proofs for these statements in Section A.2.

Stochastic IHT for Functions with Restricted Smoothness and Strong Convexity. We also provide a streamlined analysis for stochastic IHT under standard assumptions. Compared to [JTK14] which achieves similar guarantees (in particular both suffer from a blow-up in sparsity that is quadratic in the restricted condition number of f) we significantly simplify the analysis and provide guarantees when only stochastic gradients are available. Notably, the quadratic dependence in condition number can be improved to linear, at the expense of a significantly more sophisticated algorithm which requires a longer running time [AS20].

Theorem 2. Let $f : \mathbb{R}^N \rightarrow \mathbb{R}$ be a function with a k^* -sparse minimizer θ^* . Let $\beta > \alpha > 0$ be parameters, let $k = C \cdot k^* \cdot (\beta/\alpha)^2$ for some appropriately chosen constant C , and suppose that f is $(2k + k^*, \beta)$ -smooth and $(k + k^*, \alpha)$ -strongly convex in the sense that

$$f(\theta + \delta) \geq f(\theta) + \nabla f(\theta)^\top \delta + \frac{\alpha}{2} \|\delta\|^2, \text{ for all } \theta, \delta \text{ s. t. } \|\delta\|_0 \leq k + k^*.$$

For initial parameters θ_0 and precision $\epsilon > 0$, given access to stochastic gradients with variance σ , IHT (3) converges in $O\left(\frac{\beta}{\alpha} \cdot \ln \frac{\|\theta_0 - \theta^*\|}{\epsilon/\beta + \sigma^2/(\alpha\beta)}\right)$ iterations to a point θ with $\|\theta\|_0 \leq k$, such that

$$\mathbb{E} \left[\|\theta - \theta^*\|^2 \right] \leq \frac{\epsilon}{\beta} + \frac{2\sigma^2}{\alpha\beta}$$

and

$$\mathbb{E} [f(\theta) - f(\theta^*)] \leq \epsilon + \frac{2\sigma^2}{\alpha}.$$

We prove this statement in Section A.3.

Finding a Sparse Nearly-Stationary Point. We can further relax the assumptions, and show that IHT can recover sparse iterates that are nearly-stationary (i.e. have small gradient norm). Finding nearly-stationary points is a standard objective in non-convex optimization. However, enforcing the sparsity guarantee is not. Here we show that IHT can further provide stationarity guarantees even with minimal assumptions.

Intuitively, this results suggests that error feedback may not be necessary for converging to stationary points under parameter sparsity.

In this case, we alternate IHT steps with optimization steps over the sparse support, which reduce the norm of the gradient restricted to the support to some target error ϵ .

Theorem 3. Let $f : \mathbb{R}^N \rightarrow \mathbb{R}$. Let $\beta, \epsilon > 0$ be parameters, let k be the target sparsity, and suppose that f is $(2k, \beta)$ -smooth. Furthermore, suppose that after each IHT step with a step size $\eta = \beta^{-1}$ and target sparsity k , followed by optimizing over the support to gradient norm ϵ , all the obtained iterates satisfy

$$\|\theta\|_\infty \leq R_\infty(k).$$

For initial parameters θ_0 and precision $\epsilon > 0$, given access to stochastic gradients with variance σ , in $O\left(\beta(f(\theta_0) - f(\theta_T)) \cdot \min\left\{\frac{1}{\epsilon^2}, \frac{1}{\beta^2 k R_\infty(\eta, k)^2}, \frac{1}{\beta^2 \sigma^2}\right\}\right)$ iterations we can obtain an k -sparse iterate θ such that

$$\|\nabla f(\theta)\|_\infty = O\left(\beta R_\infty(k) \sqrt{k} + \beta \sigma + \epsilon\right).$$

We provide proofs in Section A.4.

Proof Approach. Let us briefly explain the intuition behind our theoretical analyses. We can view IHT as a version of projected gradient descent, where iterates are projected onto the *non-convex* domain of k -sparse vectors.

In general, when the domain is convex, projections do not hurt convergence. This is because under specific assumptions, gradient descent makes progress by provably decreasing the ℓ_2 distance between the iterate and the optimal solution. As projecting the iterate back onto the convex domain can only improve the distance to the optimum, convergence is unaffected even in those constrained settings.

In our setting, as the set of k -sparse vectors is non-convex, the distance to the optimal solution can increase. However, we can show a trade-off between how much this distance increases and the ratio between the target and optimal sparsity k/k^* .

Intuition. Intuitively, consider a point $\tilde{\theta}$ obtained by taking a gradient step

$$\tilde{\theta} = \theta - \eta \nabla f(\theta). \quad (6)$$

While this step provably decreases the distance to the optimum $\|\tilde{\theta} - \theta^*\| \ll \|\theta - \theta^*\|$, after applying the projection by moving to the projection $T_k(\tilde{\theta})$, the distance $\|T_k(\tilde{\theta}) - \theta^*\|$ may increase again. The key is that this increase can be controlled. For example, the new distance can be bounded via triangle inequality by

$$\|T_k(\tilde{\theta}) - \theta^*\| \leq \|\tilde{\theta} - T_k(\tilde{\theta})\| + \|\tilde{\theta} - \theta^*\| \leq 2\|\tilde{\theta} - \theta^*\|.$$

The last inequality follows from the fact that by definition $T_k(\tilde{\theta})$ is the closest k -sparse vector to $\tilde{\theta}$ in ℓ_2 norm, and thus the additional distance paid to move to this projected point is bounded by $\|\tilde{\theta} - \theta^*\|$. Thus, if the gradient step (6) made sufficient progress, for example $\|\tilde{\theta} - \theta^*\| \leq \frac{1}{3}\|\theta - \theta^*\|$ we can conclude that the additional truncating step does not fully undo progress, as

$$\|T_k(\tilde{\theta}) - \theta^*\| \leq 2\|\tilde{\theta} - \theta^*\| \leq \frac{2}{3}\|\theta - \theta^*\|,$$

so the iterate still converges to the optimum.

In reality, we can not always guarantee that a single gradient step reduces the distance by a large constant fraction – as a matter of fact this is determined by how well the function f is conditioned. However we can reduce the lost progress that is caused by the truncation step simply by increasing the number of nonzeros of the target solution, i.e. increasing the ratio k/k^* .

This is captured by the following crucial lemma which also appears in [JTK14]. A short proof can be found in Section A.5.

Lemma 1. *Let θ^* be a k^* -sparse vector, and let θ be an n -sparse vector. Then*

$$\frac{\|T_k(\theta) - \theta\|^2}{n - k} \leq \frac{\|\theta^* - \theta\|^2}{n - k^*}.$$

Its usefulness is made obvious in the proof of Theorem 2 which explicitly tracks as a measure of progress the distance between the current iterate and the sparse optimum. This is shown in detail in Section A.3.

The proofs of Theorems 1 and 3 are slightly more complicated, as they track progress in terms of function value rather than distance to the sparse optimum. However, similar arguments based on Lemma 1 still go through. An added benefit is that these analyses show that alternating IHT steps with gradient steps over the sparse support can only help convergence, as these additional steps further decrease error in function value. This fact is important to theoretically justify the performance of the AC/DC algorithm. In the following sections we provide proofs for our theorems.

A.2 Stochastic IHT for Non-Convex Functions with Concentrated PL Condition

In this section we show that the ideas developed before apply to non-convex settings. We analyze IHT for a class of functions that satisfy a special version of the Polyak-Łojasiewicz (PL) condition [KNS16] which is standard in non-convex optimization, and certain versions of it were essential in several works analyzing the convergence of training methods for deep neural networks [LZB20, AZLS19]. Usually this condition says that small gradient norm i.e. approximate stationarity implies closeness to optimum in function value. Here we use the stronger (r, α) -CPL condition (see Equation 5), which considers the norm of the gradient contributed by its largest coordinates in absolute value.

We prove strong convergence bounds for functions that satisfy the CPL condition. Compared to the classical Polyak-Łojasiewicz condition, this adds the additional assumption that most of the mass of the gradient is concentrated on a small subset of coordinates. This phenomenon has been witnessed in several instances, and is implicitly used in [LSB⁺19].

Before proceeding with the proof we again provide a few useful lemmas.

Lemma 2. *If f is (ℓ, β) -smooth, then for any ℓ -sparse δ , one has that*

$$f(\theta + \delta) \leq f(\theta) + \frac{\beta}{2} \left\| \left(\frac{1}{\beta} \nabla f(\theta) + \delta \right)_{\text{supp}(\delta)} \right\|^2 - \frac{1}{2\beta} \left\| \nabla f(\theta)_{\text{supp}(\delta)} \right\|^2.$$

Proof. Applying smoothness we bound

$$\begin{aligned} f(\theta + \delta) &\leq f(\theta) + \langle \nabla f(\theta), \delta \rangle + \frac{\beta}{2} \|\delta\|^2 \\ &= f(\theta) + \frac{1}{2\beta} \|\nabla f(\theta)\|^2 + \langle \nabla f(\theta), \delta \rangle + \frac{\beta}{2} \|\delta\|^2 - \frac{1}{2\beta} \|\nabla f(\theta)\|^2 \\ &= f(\theta) + \frac{1}{2} \left\| \frac{1}{\sqrt{\beta}} \nabla f(\theta) + \sqrt{\beta} \delta \right\|^2 - \frac{1}{2\beta} \|\nabla f(\theta)\|^2 \\ &= f(\theta) + \frac{\beta}{2} \left\| \frac{1}{\beta} \nabla f(\theta) + \delta \right\|^2 - \frac{1}{2\beta} \|\nabla f(\theta)\|^2. \end{aligned}$$

Next we notice that the contributions of the two terms $\frac{\beta}{2} \left\| \frac{1}{\beta} \nabla f(\theta) + \delta \right\|^2$ and $\frac{1}{2\beta} \|\nabla f(\theta)\|^2$ exactly match on the coordinates not touched by δ . Hence everything outside the support of δ cancels out, which yields the desired conclusion. \square

We require another lemma which will be very useful in the analysis.

Lemma 3. Let $\theta, g \in \mathbb{R}^N$ such that $\text{supp}(\theta) = S$, and let S', S^* be some arbitrary subsets, with $|S'| = |S| > |S^*|$. Furthermore suppose that

$$T_k(\theta + g) = (\theta + g)_{S'} .$$

Then

$$\left\| (\theta + g)_{S \setminus S'} \right\|^2 - \|g_{S \cup S'}\|^2 \leq \left\| (\theta + g)_{Z \setminus S'} \right\|^2 - \|g_{S^*}\|^2 ,$$

for some set Z such that $|Z \setminus S'| \leq 2|S^*|$.

Proof. We prove this as follows. We write

$$\begin{aligned} \left\| (\theta + g)_{S \setminus S'} \right\|^2 - \|g_{S \cup S'}\|^2 &= \left\| (\theta + g)_{(S^* \cup S) \setminus S'} \right\|^2 - \left\| (\theta + g)_{S^* \setminus (S \cup S')} \right\|^2 - \|g_{S \cup S'}\|^2 \\ &= \left\| (\theta + g)_{(S^* \cup S) \setminus S'} \right\|^2 - \|g_{S^* \setminus (S \cup S')} \|^2 - \|g_{S \cup S'}\|^2 \\ &= \left\| (\theta + g)_{(S^* \cup S) \setminus S'} \right\|^2 - \|g_{S^* \cup S \cup S'}\|^2 \\ &= \left\| (\theta + g)_{(S^* \cup S) \setminus S'} \right\|^2 - \|g_{S' \setminus (S^* \cup S)}\|^2 - \|g_{S^* \cup S}\|^2 . \end{aligned}$$

Since

$$\|g_{S' \setminus (S^* \cup S)}\|^2 = \left\| (\theta + g)_{S' \setminus (S^* \cup S)} \right\|^2 \geq \|(\theta + g)_R\|^2 ,$$

where R is a subset $R \subseteq (S^* \cup S) \setminus S'$ with $|R| = |S' \setminus (S^* \cup S)|$. Such a set definitely exists as

$$|(S^* \cup S) \setminus S'| \geq |S \setminus S'| = |S' \setminus S| \geq |S' \setminus (S^* \cup S)| = |R| .$$

Hence we obtain that

$$\left\| (\theta + g)_{S \setminus S'} \right\|^2 - \|g_{S \cup S'}\|^2 \leq \left\| (\theta + g)_{((S^* \cup S) \setminus S') \setminus R} \right\|^2 - \|g_{S^* \cup S}\|^2 .$$

Note that

$$\begin{aligned} |((S^* \cup S) \setminus S') \setminus R| &= |(S^* \cup S) \setminus S'| - |R| = |(S^* \cup S) \setminus S'| - |S' \setminus (S^* \cup S)| \\ &\leq (|S^*| + |S \setminus S'|) - (|S' \setminus S| - |S^*|) \\ &= 2|S^*| . \end{aligned}$$

This concludes the proof. \square

Using this we derive the following useful corollary.

Corollary 2. Let $\theta, g \in \mathbb{R}^N$ such that $\text{supp}(\theta) = S$, and let S', S^* be arbitrary subsets, with $|S'| = |S| > |S^*|$. Furthermore suppose that

$$T_k(\theta + g) = (\theta + g)_{S'} .$$

Then one has that

$$\left\| (\theta + g)_{S \setminus S'} \right\|^2 - \|g_{S \cup S'}\|^2 \leq \frac{2|S^*| + |\text{supp}(\theta^*)|}{|S'| - |\text{supp}(\theta^*)|} \cdot \|(\theta + g)_T - \theta^*\|^2 - \|g_{S^*}\|^2 ,$$

for some T , such that $|T| \leq 2|S^*| + |\text{supp}(\theta^*)| + |S'|$ and $\text{supp}(\theta^*) \subseteq T$.

Proof. Using Lemma 3 we can write

$$\left\| (\theta + g)_{S \setminus S'} \right\|^2 - \|g_{S \cup S'}\|^2 \leq \left\| (\theta + g)_{Z \setminus S'} \right\|^2 - \|g_{S^*}\|^2 \leq \left\| (\theta + g)_{(Z \cup \text{supp}(\theta^*)) \setminus S'} \right\|^2 - \|g_{S^*}\|^2$$

$$= \left\| (\theta + g)_{Z \cup \text{supp}(\theta^*) \cup S'} - (\theta + g)_{S'} \right\|^2 - \|g_{S^*}\|^2 .$$

where $|Z \cup \text{supp}(\theta^*) \cup S'| \leq 2|S^*| + |\text{supp}(\theta^*)| + |S'|$. Applying Lemma 1 we furthermore obtain that

$$\begin{aligned} \|(\theta + g)_{Z \cup S^* \cup S'} - (\theta + g)_{S'}\|^2 &\leq \frac{|Z \cup \text{supp}(\theta^*) \cup S'| - |S'|}{|Z \cup \text{supp}(\theta^*) \cup S'| - |\text{supp}(\theta^*)|} \cdot \left\| (\theta + g)_{Z \cup \text{supp}(\theta^*) \cup S'} - \theta^* \right\|^2 \\ &\leq \frac{2|S^*| + |\text{supp}(\theta^*)|}{|S'| - |\text{supp}(\theta^*)|} \cdot \left\| (\theta + g)_{Z \cup \text{supp}(\theta^*) \cup S'} - \theta^* \right\|^2 . \end{aligned}$$

□

We can now proceed with the main proof.

Proof of Theorem 1. For simplicity we first provide the proof for the deterministic version, which roughly follows the ideas described in [JTK14]. Afterwards, we extend it to the stochastic setting. To simplify notation, throughout this proof we will use $\eta = \frac{1}{\beta}$.

Using Lemma 2 we can write that for the update

$$\delta = T_k(\theta - \eta \nabla f(\theta)) - \theta ,$$

we have

$$\begin{aligned} f(\theta') &\leq f(\theta) + \frac{1}{2\eta} \left\| \left(\frac{1}{\beta} \nabla f(\theta) + T_k(\theta - \eta \nabla f(\theta)) - \theta \right)_{\text{supp}(\delta)} \right\|^2 - \frac{1}{2\eta} \left\| \eta \nabla f(\theta)_{\text{supp}(\delta)} \right\|^2 \\ &= f(\theta) + \frac{1}{2\eta} \left\| (T_k(\theta - \eta \nabla f(\theta)) - (\theta - \eta \nabla f(\theta)))_{\text{supp}(\delta)} \right\|^2 - \frac{1}{2\eta} \left\| \eta \nabla f(\theta)_{\text{supp}(\delta)} \right\|^2 . \end{aligned}$$

At this point we use the fact that by definition $\text{supp}(\delta) = S' \cup S$. Furthermore we see that

$$\|(T_k(\theta - \eta \nabla f(\theta)) - (\theta - \eta \nabla f(\theta)))_{S' \cup S}\|^2 = \|(\theta - \eta \nabla f(\theta))_{S \setminus S'}\|^2$$

since $T_k(\theta - \eta \nabla f(\theta))$ exactly matches $\theta - \eta \nabla f(\theta)$ for the coordinates in S' , and is 0 for all the others. Thus we have that

$$f(\theta') \leq f(\theta) + \frac{1}{2\eta} \left(\left\| (\theta - \eta \nabla f(\theta))_{S \setminus S'} \right\|^2 - \left\| \eta \nabla f(\theta)_{S \cup S'} \right\|^2 \right) .$$

In order to apply the CPL property, we need to relate this to the contribution to the gradient norm given by the heavy signal $\|T_{k^*}(\nabla f(\theta))\|^2$. Let S^* be the support of $T_{k^*}(\nabla f(\theta))$. We apply Corollary 2 to further bound

$$f(\theta') \leq f(\theta) + \frac{1}{2\eta} \left(\frac{3k^*}{k - k^*} \|(\theta - \eta \nabla f(\theta))_T - \theta^*\|^2 - \left\| \eta \nabla f(\theta)_{S^*} \right\|^2 \right) ,$$

where $|T| \leq 3k^* + k$. Now we apply the CPL property as follows. From Lemma 6 we upper bound

$$\begin{aligned} \|(\theta - \eta \nabla f(\theta))_T - \theta^*\|^2 &\leq \|(\theta - \eta \nabla f(\theta))_{T \cup S} - \theta^*\|^2 \\ &\leq \frac{8}{\alpha} (f(\theta - \eta \nabla f(\theta))_{T \cup S} - f(\theta^*)) \\ &\leq \frac{8}{\alpha} (f(\theta) - f(\theta^*)) , \end{aligned}$$

In the first inequality we used the fact that $\text{supp}(\theta^*) \subseteq T$. In the second one we applied Lemma 6. In the third one we applied $(3k^* + 2k, \beta)$ -smoothness (since $|T \cup S| \leq 3k^* + 2k$) together with the fact that $\eta = 1/\beta$, and so $f(\theta - \eta \nabla f(\theta))_{T \cup S} \leq f(\theta)$.

Similarly we apply the CPL inequality to conclude that

$$f(\theta') \leq f(\theta) + \frac{1}{2\eta} \left(\frac{3k^*}{k-k^*} \cdot \frac{8}{\alpha} (f(\theta) - f(\theta^*)) - \eta^2 \cdot \frac{\alpha}{2} (f(\theta) - f(\theta^*)) \right).$$

Thus equivalently:

$$\begin{aligned} f(\theta') - f(\theta^*) &\leq (f(\theta) - f(\theta^*)) \left(1 + \frac{12k^*}{k-k^*} \cdot \frac{1}{\eta\alpha} - \frac{\eta\alpha}{4} \right) \\ &= (f(\theta) - f(\theta^*)) \left(1 + \frac{12k^*}{k-k^*} \cdot \kappa - \frac{1}{4\kappa} \right), \end{aligned}$$

where $\kappa = \beta/\alpha = 1/(\eta\alpha)$. Since $k \geq k^* \cdot (96\kappa^2 + 1)$ we equivalently have that $\frac{k^*}{k-k^*} \leq \frac{1}{96\kappa^2}$ and thus $\frac{12k^*}{k-k^*} \cdot \kappa \leq \frac{1}{8\kappa}$. Hence

$$f(\theta') - f(\theta^*) \leq (f(\theta) - f(\theta^*)) \cdot \left(1 - \frac{1}{8\kappa} \right).$$

This shows that in $O\left(\kappa \ln \frac{f(\theta_0) - f(\theta^*)}{\epsilon}\right)$ we reach a point θ such that $f(\theta) - f(\theta^*) \leq \epsilon$, which concludes the proof.

We extend this proof to the stochastic version in Section A.5. □

Now let us show how Corollary 1 follows from the same analysis.

Proof of Corollary 1. The proof extends from that to Theorem 1. The difference we need to handle is the error introduced by performing Δ_c passes through the data instead of a single stochastic gradient step. Suppose that before starting the dense training phase, the current iterate is θ . The key is to bound the error introduced by performing these Δ_c passes instead of simply changing the iterate by $-\eta \nabla f(\theta)$, prior to applying the pruning step. To do so we first upper bound the change in the iterate after Δ_c passes, which lead to a new iterate $\tilde{\theta}$. Since we perform $\Delta_c \cdot B$ iterations instead of a single one, we damp down our step size by setting $\eta' = \eta/(\Delta_c B)$, and measure the movement we made compared to the one we would have made with a single deterministic gradient step.

Using the Lipschitz property of the functions in the decomposition, we see that each step changes our iterate by at most $\eta' L$ in ℓ_2 norm. Hence over Δ_c passes through the data, each of which involves B mini-batches, the total change in the iterate is at most:

$$\|\tilde{\theta} - \theta\| \leq \Delta_c \cdot B \cdot \eta' L = \eta L.$$

Using the smoothness property of each f_i this guarantees that for all the seen iterates $\hat{\theta}$, the gradients of f_i 's never deviate significantly from their values at the original point θ :

$$\|\nabla f_i(\hat{\theta}) - \nabla f_i(\theta)\| \leq \beta \eta L,$$

Thus if we interpret the scaled total movement in iterate $\frac{1}{\eta}(\theta - \tilde{\theta})$ as a gradient mapping, it satisfies

$$\left\| \frac{1}{\eta}(\theta - \tilde{\theta}) - \nabla f(\theta) \right\| \leq \beta \eta L.$$

Applying the analysis for the stochastic version of Theorem 1 we can treat the error in the gradient mapping exactly as the stochastic noise. For the specific choice of step size used there $\eta = 1/\Theta(\beta)$, we thus get that

$$\left\| \frac{1}{\eta}(\theta - \tilde{\theta}) - \nabla f(\theta) \right\| = O(L),$$

thus the Δ_c passes over the data can be treated as a single gradient step with error $O(L)$, which introduces in total

$$\eta \|\nabla f(\theta)\| \cdot O(L) + \frac{\beta}{2} \eta^2 L^2 = O(\eta L^2),$$

error in the function value decrease per iteration with step size η . Treating the $O(L^2)$ term the same as stochastic noise σ^2 in the original analysis, we conclude the proof.

Note that the steps performed during the sparse training phases do not affect convergence as they can only improve error in function value, which is the main quantity that our analysis tracks. \square

A.3 Stochastic IHT for Functions with Restricted Smoothness and Strong Convexity

Before proceeding with the proof for Theorem 2, we provide a few useful statements. Additionally we provide a few more useful statements related to the fact that f is well conditioned along sparse directions.

Lemma 4. *If f is $(2k + k^*, \beta)$ -smooth and $(k + k^*, \alpha)$ -strongly convex then*

$$\begin{aligned} f(\theta^*) &\geq f(\theta) + \langle \nabla f(\theta)_{S \cup S' \cup S^*}, \theta^* - \theta \rangle + \frac{\alpha}{2} \|\theta - \theta^*\|^2, \\ f\left(\theta - \frac{1}{\beta} \nabla f(\theta)_{S \cup S' \cup S^*}\right) &\leq f(\theta) - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S' \cup S^*}\|^2. \end{aligned}$$

Proof. The former follows directly from the definition. For the latter we have

$$\begin{aligned} f\left(\theta - \frac{1}{\beta} \nabla f(\theta)_{S \cup S' \cup S^*}\right) &\leq f(\theta) - \frac{1}{\beta} \langle \nabla f(\theta), \nabla f(\theta)_{S \cup S' \cup S^*} \rangle + \frac{\beta}{2} \|\nabla f(\theta)_{S \cup S' \cup S^*}\|^2 \\ &= f(\theta) - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S' \cup S^*}\|^2. \end{aligned}$$

\square

Finally we can prove the main statement:

Proof. We will track progress by measuring the distance $\|\theta - \theta^*\|^2$. To do so we write

$$\begin{aligned} \|\theta' - \theta^*\|^2 &= \|T_k(\theta - \eta g_\theta) - \theta^*\|^2 \\ &= \|(\theta - \eta g_\theta - \theta^*) + (T_k(\theta - \eta g_\theta) - (\theta - \eta g_\theta))\|^2 \\ &= \|(\theta - \eta g_\theta - \theta^*)_{S' \cup S^*} + (T_k((\theta - \eta g_\theta)_{S' \cup S^*}) - (\theta - \eta g_\theta)_{S' \cup S^*})\|^2. \end{aligned}$$

The last identity follows from the fact that the term inside the norm is only supported at $S' \cup S^*$. Thus, after applying the triangle inequality, we obtain that:

$$\|\theta' - \theta^*\|^2 \leq (\|(\theta - \eta g_\theta)_{S' \cup S^*} - \theta^*\| + \|T_k((\theta - \eta g_\theta)_{S' \cup S^*}) - (\theta - \eta g_\theta)_{S' \cup S^*}\|)^2.$$

The second term can now be bounded using Lemma 1. Since the sparsity of the projected point is $k + k^*$ we have that

$$\begin{aligned} \|T_k((\theta - \eta g_\theta)_{S' \cup S^*}) - (\theta - \eta g_\theta)_{S' \cup S^*}\|^2 &\leq \frac{k + k^* - k}{k + k^* - k^*} \|(\theta - \eta g_\theta)_{S' \cup S^*} - \theta^*\|^2 \\ &= \frac{k^*}{k} \cdot \|(\theta - \eta g_\theta)_{S' \cup S^*} - \theta^*\|^2. \end{aligned}$$

Therefore:

$$\|\theta' - \theta^*\|^2 \leq \left(1 + \sqrt{\frac{s^*}{s}}\right)^2 \cdot \|(\theta - \eta g_\theta)_{S' \cup S^*} - \theta^*\|^2 \leq \left(1 + \sqrt{\frac{s^*}{s}}\right)^2 \cdot \|(\theta - \eta g_\theta)_{S \cup S' \cup S^*} - \theta^*\|^2 .$$

The second inequality follows from the fact that as we increase the support of $\theta - \eta g_\theta$ to also include $S \setminus (S' \cup S^*)$, the norm inside can only increase. Finally by expanding, we write:

$$\begin{aligned} \|(\theta - \eta g_\theta)_{S \cup S' \cup S^*} - \theta^*\|^2 &= \|\theta - \eta(g_\theta)_{S \cup S' \cup S^*} - \theta^*\|^2 \\ &= \|\theta - \theta^*\|^2 - 2\eta \langle (g_\theta)_{S \cup S' \cup S^*}, \theta - \theta^* \rangle + \eta^2 \|(g_\theta)_{S \cup S' \cup S^*}\|^2 \\ &\leq \|\theta - \theta^*\|^2 - 2\eta \langle \nabla f(\theta)_{S \cup S' \cup S^*}, \theta - \theta^* \rangle + 2\eta^2 \|\nabla f(\theta)_{S \cup S' \cup S^*}\|^2 \\ &\quad + \underbrace{2\eta \langle (\nabla f(\theta) - g_\theta)_{S \cup S' \cup S^*}, \theta - \theta^* \rangle + 2\eta^2 \|(\nabla f(\theta) - g_\theta)_{S \cup S' \cup S^*}\|^2}_{\zeta} , \end{aligned}$$

where we use ζ to denote the error term introduced by using stochastic gradients. Next we bound the first part of the term above. To do so we use lemma to write

$$\begin{aligned} &\|\theta - \theta^*\|^2 - 2\eta \langle \nabla f(\theta)_{S \cup S' \cup S^*}, \theta - \theta^* \rangle + 2\eta^2 \|\nabla f(\theta)_{S \cup S' \cup S^*}\|^2 \\ &\leq \|\theta - \theta^*\|^2 - 2\eta \left(f(\theta) - f(\theta^*) - \frac{\alpha}{2} \|\theta - \theta^*\|^2 \right) + 2\eta^2 \|\nabla f(\theta)_{S \cup S' \cup S^*}\|^2 \\ &= \|\theta - \theta^*\|^2 (1 - \eta\alpha) - 2\eta \left(f(\theta) - f(\theta^*) - \eta \|\nabla f(\theta)_{S \cup S' \cup S^*}\|^2 \right) \\ &= \|\theta - \theta^*\|^2 (1 - \eta\alpha) - 2\eta \left(f(\theta) - f(\theta^*) - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S' \cup S^*}\|^2 \right) - 2\eta \cdot \left(\frac{1}{2\beta} - \eta \right) \|\nabla f(\theta)_{S \cup S' \cup S^*}\|^2 \\ &\leq \|\theta - \theta^*\|^2 (1 - \eta\alpha) - 2\eta \cdot \left(\frac{1}{2\beta} - \eta \right) \|\nabla f(\theta)_{S \cup S' \cup S^*}\|^2 . \end{aligned}$$

For the first inequality we used the restricted strong convexity property, while for the second we used the restricted smoothness property. Now for the error term we have that:

$$\mathbb{E}[\zeta|\theta] \leq 2\eta^2 \sigma^2 ,$$

which enables us to conclude that

$$\mathbb{E} \left[\|\theta' - \theta^*\|^2 \middle| \theta \right] \leq \left(1 + \sqrt{\frac{k^*}{k}}\right)^2 \cdot \left(\|\theta - \theta^*\|^2 (1 - \eta\alpha) - 2\eta \cdot \left(\frac{1}{2\beta} - \eta \right) \|\nabla f(\theta)_{S \cup S' \cup S^*}\|^2 + 2\eta^2 \sigma^2 \right) .$$

Setting $\eta = \frac{1}{2\beta}$ this gives us that

$$\mathbb{E} \left[\|\theta' - \theta^*\|^2 \middle| \theta \right] \leq \left(1 + \sqrt{\frac{s^*}{s}}\right)^2 \cdot \left(\|\theta - \theta^*\|^2 \left(1 - \frac{1}{2} \cdot \frac{\alpha}{\beta}\right) + \frac{1}{2} \left(\frac{\sigma}{\beta}\right)^2 \right) .$$

Thus for as long as $\|\theta - \theta^*\|^2 \geq 2\frac{\sigma^2}{\alpha\beta}$ we have that

$$\mathbb{E} \left[\|\theta' - \theta^*\|^2 \middle| \theta \right] \leq \left(1 + \sqrt{\frac{k^*}{k}}\right)^2 \cdot \left(1 - \frac{1}{4} \cdot \frac{\alpha}{\beta}\right) \cdot \|\theta - \theta^*\|^2 .$$

We see that the expected squared distance contracts by setting the ratio k/k^* sufficiently large. Indeed if $k \geq 81(\beta/\alpha)^2 \cdot k^*$ we have that:

$$\mathbb{E} \left[\|\theta' - \theta^*\|^2 \middle| \theta \right] \leq \left(1 + \frac{1}{9} \cdot \frac{\alpha}{\beta}\right)^2 \cdot \left(1 - \frac{1}{4} \cdot \frac{\alpha}{\beta}\right) \cdot \|\theta - \theta^*\|^2 \leq \left(1 - \frac{1}{36} \cdot \frac{\alpha}{\beta}\right) \cdot \|\theta - \theta^*\|^2 .$$

Taking expectation over the entire history of iterates we can thus conclude that after $T = O\left(\frac{\beta}{\alpha} \cdot \ln \frac{\|\theta_0 - \theta^*\|^2}{\epsilon/\beta + \sigma^2/\alpha\beta}\right)$ iterations we obtain

$$\mathbb{E} \left[\|\theta_T - \theta^*\|^2 \right] \leq \frac{\epsilon}{\beta} + \frac{2\sigma^2}{\alpha\beta} .$$

Applying the restricted smoothness property, this also gives us that:

$$\mathbb{E} [f(\theta_T) - f(\theta^*)] \leq \epsilon + \frac{2\sigma^2}{\alpha} .$$

which is what we wanted. \square

A.4 Finding a Sparse Nearly-Stationary Point

For simplicity, we first prove the deterministic version of the theorem, where $\sigma = 0$. We show how to extend this proof to the stochastic version in Section A.5.

Proof. The proof is similar to that of Theorem 1, but in addition requires that the algorithm alternates standard IHT steps with optimizing inside the support of the iterate. More precisely given a current iterate supported at S , we additionally run an inner loop which optimizes only over the coordinate in S , seeking a near stationary point θ such that $\|\nabla f(\theta)_S\| \leq \epsilon$. Thus in our analysis we can assume that before performing a gradient, followed by a pruning step, our current iterate supported at S satisfies $\|\nabla f(\theta)_S\| \leq \epsilon$. Hence following the previous analyses and setting $\eta = 1/\beta$, we have:

$$\begin{aligned} f(\theta') &\leq f(\theta) + \frac{1}{2\eta} \left(\left\| (\theta - \eta \nabla f(\theta))_{S \setminus S'} \right\|^2 - \left\| \eta \nabla f(\theta)_{S \cup S'} \right\|^2 \right) \\ &\leq f(\theta) + \frac{1}{2\eta} \left(2 \left\| \theta_{S \setminus S'} \right\|^2 + 2\eta^2 \epsilon^2 - \left\| \eta \nabla f(\theta) \right\|_\infty^2 \right) . \end{aligned}$$

For the second inequality we first applied triangle inequality together with the near stationarity condition for θ to bound

$$\left\| (\theta - \eta \nabla f(\theta))_{S \setminus S'} \right\| \leq \left\| \theta_{S \setminus S'} \right\| + \left\| \eta \nabla f(\theta)_{S \setminus S'} \right\| \leq \left\| \theta_{S \setminus S'} \right\| + \eta \epsilon ,$$

then applied $(a+b)^2 \leq 2a^2 + 2b^2$. In addition we used the fact that

$$\left\| \nabla f(\theta)_{S \cup S'} \right\|_\infty = \left\| \nabla f(\theta) \right\|_\infty .$$

This follows from a simple case analysis. If one of the coordinates of the gradient with the largest absolute value lies in $S \cup S'$, we are done. Otherwise, we have two possibilities. Either S' is different from S , so $S \cup S'$ contains one coordinate outside of S . Since these coordinates are obtained by hard thresholding and θ is supported only at S , the absolute value of $\nabla f(\theta)$ at the largest coordinate in $S' \setminus S$ must be at least as large as the largest in $S^* \setminus S$, which yields our claim. Otherwise we have that $S = S'$, which means that the pruning step did not change the support, and thus

$$\left\| \eta \nabla f(\theta)_{\bar{S}} \right\|_\infty \leq \min_{i \in S} |\theta - \eta \nabla f(\theta)|_i \leq R_\infty(k) + \eta \epsilon$$

which guarantees that

$$\begin{aligned} \left\| \nabla f(\theta) \right\|_\infty &= \max \left\{ \left\| \nabla f(\theta)_S \right\|_\infty , \left\| \nabla f(\theta)_{\bar{S}} \right\|_\infty \right\} \\ &\leq \max \left\{ \epsilon , \epsilon + \frac{R_\infty(k)}{\eta} \right\} = \epsilon + \beta R_\infty(k) , \end{aligned}$$

and so we are done.

In the former case we thus see that

$$\frac{\eta}{2} \|\nabla f(\theta)\|_\infty^2 \leq f(\theta) - f(\theta') + \left(\frac{\|\theta_{S \setminus S'}\|^2}{\eta} + \eta\epsilon^2 \right) \leq f(\theta) - f(\theta') + \left(\frac{kR_\infty(k)^2}{\eta} + \eta\epsilon^2 \right).$$

Telescoping over T iterations we see that

$$\frac{\eta}{2} \sum_{t=0}^{T-1} \|\nabla f(\theta_t)\|_\infty^2 \leq f(\theta_0) - f(\theta_T) + T \cdot \left(\frac{kR_\infty(k)^2}{\eta} + \eta\epsilon^2 \right)$$

and so returning a random point θ among those witnessed during the algorithm we have

$$\mathbb{E} \left[\|\nabla f(\theta)\|_\infty^2 \right] \leq \frac{2(f(\theta_0) - f(\theta_T))}{\eta T} + 2 \left(\frac{kR_\infty(k)^2}{\eta^2} + \epsilon^2 \right)$$

By AM-QM,

$$\mathbb{E} [\|\nabla f(\theta)\|_\infty] \leq \sqrt{\mathbb{E} [\|\nabla f(\theta)\|_\infty^2]},$$

which enables us to conclude that after sufficiently many iterations we are guaranteed to find a point such that

$$\|\nabla f(\theta)\|_\infty = O \left(\frac{R_\infty(k) \sqrt{k}}{\eta} + \epsilon \right) = O \left(\beta R_\infty(k) \sqrt{k} + \epsilon \right).$$

□

A.5 Deferred Proofs

Proof of Lemma 1. Our proofs crucially rely on the following lemma. Intuitively it shows that projecting a vector θ onto the non-convex set of sparse vectors does not increase the distance to the optimum by too much. While this is indeed always true for projections onto convex sets, in this case we can provably show that the possible increase in distance is small.

Proof. We have that the function

$$h(k) = \frac{\|T_k(\theta) - \theta\|^2}{n - k}$$

is non-increasing. Indeed, using more nonzeros can only decrease the ratio. Thus

$$\frac{\|T_k(\theta) - \theta\|^2}{n - k} \leq \frac{\|T_{k^*}(\theta) - \theta\|^2}{n - k^*} \leq \frac{\|\theta^* - \theta\|^2}{n - k^*},$$

where the last inequality follows from the fact that among all k -sparse vectors, $T_k(\theta)$ minimizes the distance to θ . □

Proof of Theorem 1 (stochastic version). Next we extend the proof of Theorem 1 to the case when only stochastic gradients are available.

Proof. Similarly to before we can write

$$\begin{aligned} f(\theta') &\leq f(\theta) + \frac{\beta}{2} \left\| \frac{1}{\beta} \nabla f(\theta) + (T_k(\theta - \eta g_\theta) - \theta) \right\|^2 - \frac{1}{2\beta} \|\nabla f(\theta)\|^2 \\ &= f(\theta) + \frac{\beta}{2} \left\| \left(T_k(\theta - \eta g_\theta) - \left(\theta - \frac{1}{\beta} \nabla f(\theta) \right) \right) \right\|_{S \cup S'}^2 - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S'}\|^2 \end{aligned}$$

$$\begin{aligned}
&\leq f(\theta) + \frac{\beta}{2} \left(\|(T_k(\theta - \eta g_\theta) - (\theta - \eta g_\theta))_{S \cup S'}\| + \left\| \left(\eta g_\theta - \frac{1}{\beta} \nabla f(\theta) \right)_{S \cup S'} \right\|^2 - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S'}\|^2 \right) \\
&\leq f(\theta) + \beta \|(T_k(\theta - \eta g_\theta) - (\theta - \eta g_\theta))_{S \cup S'}\|^2 + \beta \left\| \left(\eta g_\theta - \frac{1}{\beta} \nabla f(\theta) \right)_{S \cup S'} \right\|^2 - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S'}\|^2 \\
&= f(\theta) + \beta \left\| (\theta - \eta g_\theta)_{S \setminus S'} \right\|^2 + \beta \left\| \left(\eta g_\theta - \frac{1}{\beta} \nabla f(\theta) \right)_{S \cup S'} \right\|^2 - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S'}\|^2.
\end{aligned}$$

Next we apply Corollary 2 to further bound

$$\begin{aligned}
\left\| (\theta - \eta g_\theta)_{S \setminus S'} \right\|^2 &\leq \frac{2|S^*| + |\text{supp}(\theta^*)|}{|S'| - |\text{supp}(\theta^*)|} \cdot \|\theta - \eta g_\theta\|_T^2 - \|\eta(g_\theta)_{S^*}\|^2 + \|\eta(g_\theta)_{S \cup S'}\|^2 \\
&= \frac{3k^*}{k - k^*} \cdot \|\theta - \eta g_\theta\|_T^2 - \|\eta(g_\theta)_{S^*}\|^2 + \|\eta(g_\theta)_{S \cup S'}\|^2,
\end{aligned}$$

where $|T| \leq 3k^* + k$. Thus

$$\begin{aligned}
f(\theta') &\leq f(\theta) + \beta \left(\frac{3k^*}{k - k^*} \cdot \|\theta - \eta g_\theta\|_T^2 - \|\eta(g_\theta)_{S^*}\|^2 + \|\eta(g_\theta)_{S \cup S'}\|^2 \right) \\
&\quad + \beta \left\| \left(\eta g_\theta - \frac{1}{\beta} \nabla f(\theta) \right)_{S \cup S'} \right\|^2 - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S'}\|^2.
\end{aligned}$$

Taking expectations, and applying Lemma 5 we see that

$$\begin{aligned}
\mathbb{E}[f(\theta') - f^* | \theta] &\leq \mathbb{E}[f(\theta) - f^* | \theta] \\
&\quad + \beta \left(\frac{3k^*}{k - k^*} \cdot \|\theta - \eta \nabla f(\theta)\|_T^2 - \|\eta \nabla f(\theta)_{S^*}\|^2 + \|\eta \nabla f(\theta)_{S \cup S'}\|^2 \right) \\
&\quad + \beta \left\| \left(\eta \nabla f(\theta) - \frac{1}{\beta} \nabla f(\theta) \right)_{S \cup S'} \right\|^2 - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S'}\|^2 + 2\beta\eta^2\sigma^2 \\
&= \mathbb{E}[f(\theta) - f^* | \theta] + \beta \left(\frac{3k^*}{k - k^*} \cdot \|\theta - \eta \nabla f(\theta)\|_T^2 - \|\eta \nabla f(\theta)_{S^*}\|^2 \right) \\
&\quad + \|\nabla f(\theta)_{S \cup S'}\|^2 \left(\beta\eta^2 + \beta \left(\eta - \frac{1}{\beta} \right)^2 - \frac{1}{2\beta} \right) + 2\beta\eta^2\sigma^2 \\
&\leq \mathbb{E}[f(\theta) - f^* | \theta] + \beta \left(\frac{3k^*}{k - k^*} \cdot \|\theta - \eta \nabla f(\theta)\|_T^2 - \|\eta \nabla f(\theta)_{S^*}\|^2 \right) + 2\beta\eta^2\sigma^2,
\end{aligned}$$

where the last inequality follows from setting $\eta = \frac{1}{2\beta}$, which makes $\beta\eta^2 + \beta \left(\eta - \frac{1}{\beta} \right)^2 - \frac{1}{2\beta} = 0$.

Finally, repeating the argument used for the deterministic proof in Section A.2, we further bound

$$\begin{aligned}
\mathbb{E}[f(\theta') - f^* | \theta] &\leq f(\theta) - f^* + \beta \left(\frac{3k^*}{k - k^*} \cdot \frac{8}{\alpha} (f(\theta) - f(\theta^*)) - \eta^2 \cdot \frac{\alpha}{2} (f(\theta) - f(\theta^*)) \right) + 2\beta\eta^2\sigma^2 \\
&= (f(\theta) - f^*) \left(1 + \frac{24k^*}{k - k^*} \cdot \frac{\beta}{\alpha} - \frac{\beta\eta^2\alpha}{2} \right) + 2\beta\eta^2\sigma^2 \\
&= (f(\theta) - f^*) \left(1 + \frac{24k^*}{k - k^*} \cdot \frac{\beta}{\alpha} - \frac{\alpha}{8\beta} \right) + \frac{\sigma^2}{2\beta}.
\end{aligned}$$

Setting $k = k^* \cdot \left(384 \left(\frac{\beta}{\alpha}\right)^2 + 1\right)$ we have $\frac{24k^*}{k-k^*} \cdot \frac{\beta}{\alpha} \leq 24 \cdot \frac{1}{384 \cdot (\beta/\alpha)^2} \cdot \frac{\beta}{\alpha} = \frac{1}{16} \cdot \frac{\alpha}{\beta}$, so

$$\mathbb{E} [f(\theta') - f^* | \theta] \leq (f(\theta) - f^*) \left(1 - \frac{\alpha}{16\beta}\right) + \frac{\sigma^2}{2\beta}.$$

Thus for as long as

$$\frac{\sigma^2}{2\beta} \leq (f(\theta) - f^*) \cdot \frac{\alpha}{32\beta} \iff f(\theta) - f^* \geq 16 \cdot \frac{\sigma^2}{\alpha}$$

one has that

$$\mathbb{E} [f(\theta') - f^* | \theta] \leq (f(\theta) - f^*) \left(1 - \frac{\alpha}{32\beta}\right).$$

Taking expectation over the entire history, this shows that after $T = O\left(\left(\frac{\beta}{\alpha}\right) \ln \frac{f(\theta_0) - f^*}{\epsilon}\right)$ iterations we obtain an iterate θ_T such that

$$\mathbb{E} [f(\theta_T) - f^*] \leq \epsilon + \frac{16\sigma^2}{\alpha},$$

which concludes the proof. \square

Proof of Theorem 3 (stochastic version). We also provide the proof for the stochastic version of Theorem 3.

Proof. We follow the steps of the proof provided in Section A.4. More precisely we write:

$$\begin{aligned} f(\theta') &\leq f(\theta) + \frac{\beta}{2} \left\| \frac{1}{\beta} \nabla f(\theta) + (T_k(\theta - \eta g_\theta) - \theta) \right\|^2 - \frac{1}{2\beta} \|\nabla f(\theta)\|^2 \\ &= f(\theta) + \frac{\beta}{2} \left\| \left(T_k(\theta - \eta g_\theta) - \left(\theta - \frac{1}{\beta} \nabla f(\theta) \right) \right) \right\|_{S \cup S'}^2 - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S'}\|^2 \\ &\leq f(\theta) + \frac{\beta}{2} \left(\|T_k(\theta - \eta g_\theta) - (\theta - \eta g_\theta)\|_{S \cup S'} + \left\| \left(\eta g_\theta - \frac{1}{\beta} \nabla f(\theta) \right) \right\|_{S \cup S'} \right)^2 - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S'}\|^2 \\ &\leq f(\theta) + \beta \|T_k(\theta - \eta g_\theta) - (\theta - \eta g_\theta)\|_{S \cup S'}^2 + \beta \left\| \left(\eta g_\theta - \frac{1}{\beta} \nabla f(\theta) \right) \right\|_{S \cup S'}^2 - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S'}\|^2 \\ &= f(\theta) + \beta \left\| (\theta - \eta g_\theta)_{S \setminus S'} \right\|^2 + \beta \left\| \left(\eta g_\theta - \frac{1}{\beta} \nabla f(\theta) \right) \right\|_{S \cup S'}^2 - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S'}\|^2, \end{aligned}$$

where we used the inequality $(a+b)^2 \leq 2a^2 + 2b^2$. Applying Lemma 5 and using the fact that $\|\nabla f(\theta)_S\| \leq \epsilon$, we see that setting $\eta = 1/\beta$ we obtain:

$$\begin{aligned} \mathbb{E} [f(\theta')] &\leq f(\theta) + \beta \left(\left\| (\theta - \eta \nabla f(\theta))_{S \setminus S'} \right\|^2 + \sigma^2 \right) \\ &\quad + \beta \left(\left\| \left(\left(\eta - \frac{1}{\beta} \right) \nabla f(\theta) \right)_{S \cup S'} \right\|^2 + \sigma^2 \right) - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S'}\|^2 \\ &\leq f(\theta) + 2\beta \left(\|\theta_{S \setminus S'}\|^2 + \frac{\epsilon^2}{\beta^2} + \sigma^2 \right) - \frac{1}{2\beta} \|\nabla f(\theta)_{S \cup S'}\|^2 \\ &\leq f(\theta) + 2\beta \left(kR_\infty (\beta^{-1}, k)^2 + \frac{\epsilon^2}{\beta^2} + \sigma^2 \right) - \frac{1}{2\beta} \|\nabla f(\theta)\|_\infty^2, \end{aligned}$$

where again we used the fact that $\|\nabla f(\theta)_{S \cup S'}\|_\infty = \|\nabla f(\theta)\|_\infty$, unless $\|\nabla f(\theta)\|_\infty \leq \frac{R_\infty(\eta, k)}{\eta} = \beta R_\infty(\beta^{-1}, k)$. Thus, telescoping over T iterations we see that

$$\frac{1}{2\beta} \sum_{t=0}^{T-1} \|\nabla f(\theta_t)\|_\infty^2 \leq f(\theta_0) - f(\theta_T) + T \cdot 2\beta \left(k R_\infty(\beta^{-1}, k)^2 + \beta^{-2} \epsilon^2 + \sigma^2 \right)$$

and so returning a random point θ among those witnessed during the algorithm we have

$$\mathbb{E} \left[\|\nabla f(\theta)\|_\infty^2 \right] \leq \frac{2\beta (f(\theta_0) - f(\theta_T))}{T} + 2 \left(\beta^2 k R_\infty(\beta^{-1}, k)^2 + \epsilon^2 + \beta^2 \sigma^2 \right)$$

By AM-QM,

$$\mathbb{E} [\|\nabla f(\theta)\|_\infty] \leq \sqrt{\mathbb{E} [\|\nabla f(\theta)\|_\infty^2]}$$

which enables us to conclude that after sufficiently many iterations we are guaranteed to find a point such that

$$\|\nabla f(\theta)\|_\infty = O \left(\beta R_\infty(\eta, k) \sqrt{k} + \beta \sigma + \epsilon \right).$$

□

Miscellaneous Proofs.

Lemma 5. *Let $\sigma > 0$ and let g_θ be a stochastic gradient satisfying standard conditions:*

$$\mathbb{E}[g_\theta | \theta] = \nabla f(\theta),$$

and

$$\mathbb{E}[\|g_\theta - \nabla f(\theta)\|^2] \leq \sigma^2.$$

Then for any vector $a \in \mathbb{R}^N$ and any subset S of coordinates:

$$\|(\nabla f(\theta) + a)_S\|^2 \leq \mathbb{E} \left[\|(g_\theta + a)_S\|^2 \middle| \theta \right] \leq \|(\nabla f(\theta) + a)_S\|^2 + \sigma^2.$$

Proof. We expand the norm under the expected value as:

$$\begin{aligned} \mathbb{E} \left[\|(g_\theta + a)_S\|^2 \middle| \theta \right] &= \mathbb{E} \left[\|(\nabla f(\theta) + a)_S + (g_\theta - \nabla f(\theta))_S\|^2 \middle| \theta \right] \\ &= \mathbb{E} \left[\|(\nabla f(\theta) + a)_S\|^2 + \|(g_\theta - \nabla f(\theta))_S\|^2 + 2 \langle (\nabla f(\theta) + a)_S, (g_\theta - \nabla f(\theta))_S \rangle \middle| \theta \right] \\ &= \|(\nabla f(\theta) + a)_S\|^2 + \mathbb{E} \left[\|(g_\theta - \nabla f(\theta))_S\|^2 \middle| \theta \right] \\ &\quad + \mathbb{E} \left[2 \langle (\nabla f(\theta) + a)_S, (g_\theta - \nabla f(\theta))_S \rangle \middle| \theta \right] \\ &\leq \|(\nabla f(\theta) + a)_S\|^2 + \sigma^2. \end{aligned}$$

From the the chain of equalities above, we can trivially lower bound the expectation by $\|(\nabla f(\theta) + a)_S\|^2$, which gives us what we needed. □

A crucial step in the proof of Theorem 1 requires upper bounding the ℓ_2 distance to the closest global optimizer by the difference in function value. In general, it is known that this is implied by the Polyak-Łojasiewicz condition, and so it automatically holds for the stronger concentrated Polyak-Łojasiewicz condition. We reproduce the proof from [KNS16] for completeness.

Lemma 6. (From Polyak-Łojasiewicz to quadratic growth) Let $f : \mathbb{R}^N \rightarrow \mathbb{R}$ be a function satisfying the Polyak-Łojasiewicz inequality

$$\|\nabla f(\theta)\|^2 \geq \frac{\alpha}{2} (f(\theta) - f^*) .$$

Then there exists a global minimizer θ^* of f such that

$$f(\theta) - f^* \geq \frac{\alpha}{8} \|\theta - \theta^*\|^2 .$$

Proof. Let $g(\theta) = \sqrt{f(\theta) - f^*}$ for which we have

$$\nabla g(\theta) = \frac{1}{2\sqrt{f(\theta) - f^*}} \nabla f(\theta) .$$

Using the PL condition we have

$$\|\nabla g(\theta)\|^2 = \frac{1}{4(f(\theta) - f^*)} \cdot \|\nabla f(\theta)\|^2 \geq \frac{1}{4(f(\theta) - f^*)} \cdot \frac{\alpha}{2} \cdot (f(\theta) - f^*) = \frac{\alpha}{8} .$$

Now starting at some θ_0 , we consider the dynamic $\dot{\theta} = -\nabla g(\theta)$. We see that this always decreases function value until it reaches some θ_T for which $\nabla g(\theta_T) = 0$ and hence by the PL inequality, θ_T is a minimizer i.e. $f(\theta_T) = f^*$. Now we can write

$$\begin{aligned} g(\theta_T) &= g(\theta_0) + \int_0^T \langle \nabla g(\theta_t), \dot{\theta}_t \rangle dt = g(\theta_0) + \int_0^T \langle \nabla g(\theta_t), -\nabla g(\theta_t) \rangle dt \\ &= g(\theta_0) - \int_0^T \|\nabla g(\theta_t)\|^2 dt . \end{aligned}$$

Thus

$$g(\theta_0) - g(\theta_T) = \int_0^T \|\nabla g(\theta_t)\|^2 dt \geq \sqrt{\frac{\alpha}{8}} \cdot \int_0^T \|\nabla g(\theta_t)\| dt = \sqrt{\frac{\alpha}{8}} \cdot \int_0^T \|\dot{\theta}_t\| dt ,$$

where we used our lower bound on the norm of $\nabla g(\theta)$. Finally, we use the fact that the last integral lower bounds the total movement of θ as it moves from θ_0 to θ_T . Thus

$$\int_0^T \|\dot{\theta}_t\| dt \geq \|\theta_0 - \theta_T\| ,$$

so

$$g(\theta_0) - g(\theta_T) \geq \sqrt{\frac{\alpha}{8}} \|\theta_0 - \theta_T\| ,$$

which enables us to conclude that

$$f(\theta_0) - f^* \geq \frac{\alpha}{8} \|\theta_0 - \theta_T\|^2 ,$$

where θ_T is some global minimizer of f . This concludes the proof. \square

B Additional Experiments

B.1 ResNet50 on ImageNet

Performance of dense models. The AC/DC method has an advantage over other pruning methods of obtaining *both* sparse and dense models. The performance of the dense baseline can be recovered after fine-tuning the resulting AC/DC dense model, for a small number of epochs. Namely, we start from the best dense baseline, which is usually obtained after 85 epochs, and replace the final compression phase of 15 epochs with regular dense training; we use the same learning rate

scheduler and keep all other training hyper-parameters the same. For 80% sparsity we recover the dense baseline accuracy completely, while for 90% we are slightly below the baseline by 0.3%. We note that for 90% sparsity, when the first and last layers are dense, our fine-tuned dense model recovers the baseline accuracy fully. The results for the dense models, together with the baseline accuracy, are presented in Table 6, where (\star) denotes that the first and last layers of the network are dense.

Table 6: AC/DC Dense ResNet50

Target Sparsity	Accuracy Dense (%)	Accuracy Finetuned (%)
0%	76.84	-
80%	73.82 ± 0.02	76.83 ± 0.07
90%	73.25 ± 0.16	76.56 ± 0.1
90% \star	73.66	76.85

Table 7: AC/DC Dense MobileNetV1


Sparsity	Accuracy Dense (%)	Accuracy Finetuned (%)
0%	71.78	-
75%	68.55 ± 0.2	71.63 ± 0.1
90%	67.47 ± 0.13	70.86 ± 0.08
90% \star	68.27	70.97

Moreover, an interesting property of AC/DC is that the resulting dense networks have a small percentage of zero-valued weights, as shown in Table 8. This is most likely caused by “dead” neurons or convolutional filters resulted after each compression phase; the corresponding weights do not get re-activated during the dense stages, as they can no longer receive gradients. This can be easily seen particularly for high sparsity (95% and 98%) where a non-trivial percentage of the weights remain inactive.


Dynamics of masks and FLOPs during training. The mask dynamics, measured by the relative change between two consecutive compression masks, have an important influence on the AC/DC training process. Namely, more changes between consecutive compression masks typically imply more exploration of the weights’ space, and faster recovery from sub-optimal pruning decision, which in turn results in more accurate sparse models. As can be seen in Figure 2b, the relative mask difference between consecutive compression phases decreases during training, but it is critical to be maintained at a non-trivial level. For completeness, we also included the evolution of the validation accuracy during AC/DC training, for all sparsity levels (please see Figure 2a); at 98% sparsity in particular, it is easiest to see that dense phases enable the exploration of better pruning masks, which ensure that the sparse model improves continuously during training.

Despite the dynamics of the compression masks, we noticed that the sparsity distribution does not change significantly. This can be observed from the number of inference FLOPs per sample, at the end of each compression phase, in Figure 4a. Interestingly, as training progresses, AC/DC also induces structured sparsity, as more neurons and convolutional filters get pruned. This was previously discussed in more detail (see Table 8), but can also be deduced from the decreasing inference FLOPs at the end of each dense phase, as shown in Figure 4b.

AC/DC with uniform pruning. As discussed, for example, in [SA20], global magnitude pruning usually performs better than its uniform counterpart. Interestingly, with global magnitude pruning later layers (which also tend to be the largest) are pruned the most. Moreover, we did not encounter convergence issues caused by entire layers being pruned, as hypothesized in some previous work [EGM⁺20, JPR⁺20]. However, one concern related to global magnitude pruning is a potential FLOP inefficiency of the resulting models; in theory, this would be a consequence of the earlier layers being pruned the least. For this reason, we performed additional experiments with AC/DC at uniform sparsity, with the first and last layers dense (as commonly used in the literature [EGM⁺20, JPR⁺20]).



(a) Test FLOPs after each sparse phase



(b) Test FLOPs after each dense phase

Figure 4: Dynamics of sparse and dense inference FLOPs for ImageNet on ResNet50, as a percentage of the dense baseline FLOPs

Table 8: Accuracy, sparsity, inference FLOPs and percentage of inactive weights for the resulting AC/DC dense models on ResNet50 (before fine-tuning, one seed).

Target Sparsity	Top-1 Accuracy (%)	Inference FLOPs	Inactive Weights (%)
80	73.8	0.98×	3.2
90	73.2	0.93×	10.5
95	72.9	0.84×	22.05
98	70.8	0.66×	49.78

Our results show that there are no significant differences compared to AC/DC with global magnitude pruning. However, keeping the first and last layers dense significantly improves the results with global magnitude pruning. These observations emphasize that AC/DC is an easy-to-use method which works reliably well with different pruning criteria. For complete results, please see Table 9.

Direct comparison with Top-KAST. As previously highlighted, Top-KAST is the closest to us, in terms of validation accuracy, out of existing sparse training methods. However, for the results reported, the authors kept the first convolutional and final fully-connected layers dense. To obtain a fair comparison, we used AC/DC on the same sparse distribution, and for 90% sparsity over the pruned layers (82.57% overall network sparsity), our results improved significantly. Namely, the best sparse model reached 75.64% validation accuracy (0.6% increase from the results in Table 2), while

Table 9: AC/DC with uniform vs global magnitude pruning on ResNet50 (one seed), where (★) denotes that the first and last layers are dense.

Sparsity Distribution	Target Sparsity(%)	Global Sparsity(%)	Top-1 Accuracy (%)	FLOPs Inference
global	90	89.8	75.14	0.19×
global★	90	82.6	75.64	0.21×
uniform★	90	82.6	75.04	0.13×
global	95	94.8	73.15	0.12×
global★	95	87.2	74.16	0.14×
uniform★	95	87.2	73.28	0.08×

Table 10: Comparison with Top-KAST when pruning all layers (ResNet50)

Method	Sparsity (%)	Backward Sparsity (%)	Sparse Top-1 Accuracy (%)
AC/DC	80	80 / 0	76.3 ± 0.1
Top-KAST	80	0	75.64
Top-KAST	80	50	74.78
Top-KAST	80	80	72.19
AC/DC	90	90 / 0	75.03 ± 0.1
Top-KAST	90	0	74.42
Top-KAST	90	50	74.09
Top-KAST	90	80	73.07

the accuracy of the best dense model was 76.85% after fine-tuning. For more details, we also provide in Table 10 the results for Top-KAST when all layers are pruned, as they were provided to us by the authors. Notice that AC/DC surpasses even Top-KAST with dense back-propagation, which has limited computational savings during training.

It is important to note, however, that because of its flexibility in choosing the gradients density, Top-KAST can theoretically obtain significantly better training speed-ups than AC/DC, the latter being constrained by its dense training phases. This allows Top-KAST to improve the accuracy of the models by increasing the number of training epochs, while still enabling (theoretical) training speed-up. We present in Table 11 another comparison between AC/DC and Top-KAST, when the training time for the latter is increased 2 or 5 times; for all results (which were provided to us by the authors), the first and last layers for Top-KAST are dense. When comparing with AC/DC with all layers pruned, Top-KAST obtains better results at 98% and 95% sparsity, with increased training epochs. However, when using the same sparse distribution as Top-KAST (not pruning the first and last layers), the results for AC/DC at 95% and 98% sparsity are significantly better than Top-KAST with increased steps. For all the results reported on AC/DC the number of training steps was fixed at 100 epochs.

We note that the results obtained with AC/DC can be improved as well with increased number of training epochs. As an example, when using the same sparsity schedule extended over 150 epochs, the best sparse model obtained with AC/DC on 90% sparsity reached 75.99% accuracy, using fewer training FLOPs compared to the original dense baseline trained on 100 epochs (namely 87%). Furthermore, when we fine-tune the dense model by replacing the final 15 epochs compression phase with dense training, we obtain a dense model with 76.95% accuracy, higher than the original dense baseline.

B.2 MobileNet on ImageNet

Performance of dense models. Similar to ResNet50, we observed that dense models obtained with AC/DC are able to recover the baseline accuracy after additional fine-tuning. We performed fine-tuning identically to the ResNet50 experiments and observe that AC/DC models obtained with a 75% target sparsity recovered the baseline accuracy, while for 90% the gap is just below 1%. We present results for the (fine-tuned) dense models in Table 7, where (\star) indicates that the first layer and depth-wise convolutions were never pruned.

Table 11: Comparison with Top-KAST with increased training steps (ResNet50). (★) indicates that the first and last layers are dense for AC/DC, while this is the case for all Top-KAST results.

Method	Sparsity (%)	Backward Sparsity (%)	Sparse Top-1 Accuracy (%)	Train FLOPs (%)	Inference FLOPs (%)
AC/DC	80	80 / 0	76.3 ± 0.1	0.65×	0.29×
Top-KAST _{1×}	80	0	75.59	0.48×	0.23×
Top-KAST _{1×}	80	60	74.59	0.29×	0.23×
Top-KAST _{2×}	80	0	76.11	0.97×	0.23×
Top-KAST _{2×}	80	60	75.29	0.58×	0.23×
AC/DC	90	90 / 0	75.03 ± 0.1	0.58×	0.19×
AC/DC★	90	90 / 0	75.64	0.6×	0.21×
AC/DC★ unif.	90	90/0	75.04	0.55×	0.13×
Top-KAST _{1×}	90	0	74.65	0.42×	0.13×
Top-KAST _{1×}	90	80	73.03	0.16×	0.13×
Top-KAST _{2×}	90	0	75.35	0.84×	0.13×
Top-KAST _{2×}	90	80	74.16	0.32×	0.13×
AC/DC	95	95 / 0	73.14 ± 0.2	0.53×	0.12×
AC/DC★	95	95 / 0	74.16	0.54×	0.14×
AC/DC★ (unif)	95	95 / 0	73.28	0.51×	0.08×
Top-KAST _{1×}	95	0	71.83	0.39×	0.08×
Top-KAST _{1×}	95	90	70.42	0.1×	0.08×
Top-KAST _{2×}	95	0	73.29	0.77×	0.08×
Top-KAST _{2×}	95	90	72.42	0.19×	0.08×
Top-KAST_{5×}	95	0	74.27	1.94×	0.08×
Top-KAST _{5×}	95	90	73.17	0.48×	0.08×
AC/DC	98	98 / 0	68.44 ± 0.09	0.46×	0.06×
AC/DC★	98	98 / 0	71.27	0.48×	0.09×
Top-KAST _{1×}	98	90	67.06	0.08×	0.05×
Top-KAST _{1×}	98	95	66.46	0.06×	0.05×
Top-KAST _{2×}	98	90	68.99	0.15×	0.05×
Top-KAST _{2×}	98	85	68.87	0.12×	0.05×

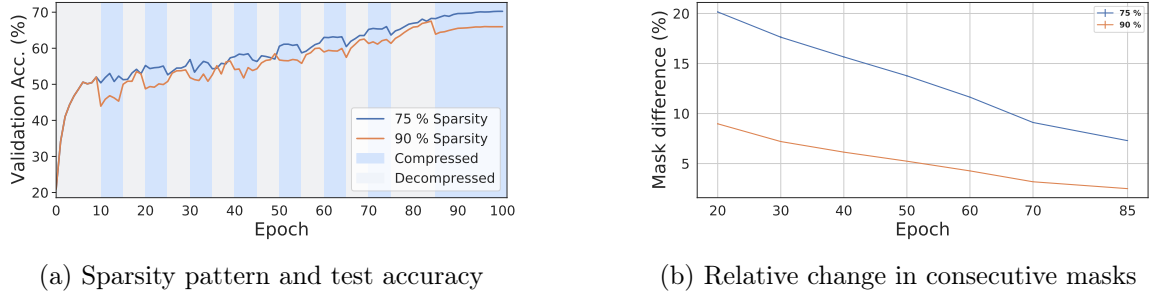


Figure 5: Validation accuracy and sparsity during training, together with differences in consecutive masks for ImageNet with MobileNetV1 using AC/DC.

Masks dynamics. Similar to ResNet50, the change between consecutive AC/DC compression masks plays an important role in obtaining accurate sparse models on MobileNet. As shown in Figure 5b, the compression masks stabilize as training progresses. For completeness, we also illustrate the evolution of the validation accuracy during AC/DC training on MobileNet, at 75% and 90% sparsity, in Figure 5a.

Comparison with RigL. We note that the results obtained by RigL [EGM⁺20] improve significantly when increasing the number of training steps 2 or 5 times. Moreover, for all results reported with RigL on MobileNet the first convolutional layer and all depth-wise convolutions are dense, whereas we do not impose such restrictions on our sparse model. Our results can further be improved by using the same sparsity distribution; namely, for 90% sparsity over the pruned parameters (88.57% overall sparsity), the best sparse model obtained with AC/DC achieved 66.56% accuracy (0.5% improvement), while the best dense improved from 67.64% to 70.97% after fine-tuning. In Table 12 we present results for AC/DC and RigL at 75% and 90% sparsity, when the latter is trained over the same number of epochs, or with 2x or 5x the number of passes through the training data. We conclude that AC/DC has very similar validation accuracy to RigL_{2x}. For 75% sparsity, AC/DC achieves similar performance with significantly fewer training and inference FLOPs than RigL. At 90% sparsity, AC/DC and RigL_{2x} are close in terms of both validation accuracy and training FLOPs; however, the validation accuracy of AC/DC can be improved by almost 0.5% when the first and depth-wise convolutional layers are kept dense. We note that RigL_{5x} has significantly higher validation accuracy, and for 75% sparsity it even matches the baseline; however, this variant of RigL also uses 2.6× and 1.5× the dense baseline training FLOPs for 75% and 90% sparsities, respectively, which makes it impractical due to its high computational training cost.

B.3 Inference Speedups

We now examine the potential for real-world speedup of models produced through our framework. For this, we use the CPU-based inference framework of [Dee21], which supports efficient inference over unstructured sparse models, and is free to use for non-commercial purposes. Specifically, we export our Pytorch-trained models to the ONNX intermediate format, preserving weight sparsity, and then execute inference on a subset of samples, at various batch sizes, measuring time per batch. We execute on an Intel i9-7980XE CPU with 16 cores and 2.60GHz core frequency. We simulate two scenarios: the first is *real-time inference*, i.e. samples are processed one at a time, in a resource-constrained environment, using only 4 cores. The second is *batch inference*, for which we pick batch size 64, in a cloud environment, for which we use all 16 cores. We measure average

Table 12: Comparison between AC/DC and RigL on MobileNet, where (\star) denotes that the first and depth-wise convolutions were kept dense.

Method	Sparsity (%)	Top-1 Accuracy (%)	Inference FLOPs	Train FLOPs
AC/DC	75	70.3	$0.35\times$	$0.65\times$
AC/DC \star	75	70.41	$0.37\times$	$0.66\times$
RigL \star (ERK)	75	68.39	$0.52\times$	$0.52\times$
RigL $\star_{2\times}$ (ERK)	75	70.49	$0.52\times$	$1.05\times$
RigL $\star_{5\times}$ (ERK)	75	71.9	$0.52\times$	$2.63\times$
AC/DC	90	66.08	$0.19\times$	$0.57\times$
AC/DC \star	90	66.56	$0.22\times$	$0.58\times$
RigL \star (ERK)	90	63.58	$0.27\times$	$0.29\times$
RigL $\star_{2\times}$ (ERK)	90	65.92	$0.27\times$	$0.59\times$
RigL $\star_{5\times}$ (ERK)	90	68.1	$0.27\times$	$1.47\times$

Table 13: Time per batch (milliseconds) using a sparse inference engine [Dee21].

Model/Setup	Real-Time Inference, 4 cores	Batch 64 Inference, 16 cores
ResNet50 ONNXRT v1.6	14.773	329.734
ResNet50 Dense	15.081	285.958
ResNet50 90% Pruned	9.46	124.193
ResNet50 90% Unif. Pruned	8.495	116.897
MobileNetV1 ONNXRT v1.6	2.552	80.748
MobileNetV1 Dense	2.513	55.845
MobileNetV1 Pruned 75%	1.96	40.976
MobileNetV1 Pruned 90%	1.468	34.909

time per batch for the sparse models against dense baselines, for which we use both the Deepsparse engine, and the ONNX runtime (ONNXRT). We present the average over 10 runs. The variance is extremely low, so we omit it for readability.

We now briefly discuss the results. First, notice that the dense baselines offer similar performance for real-time inference, but that the Deepsparse engine has a slight edge at batch 64. We will therefore compare against its timings below. The results show a speedup of 1.6x for the 90% global-pruned ResNet50 model, and 1.8x for the uniformly pruned one: the uniformly-pruned model is slightly faster, which correlates with its lower FLOP count due to the uniform pruning pattern. This pattern is preserved in MobileNetV1 experiments, although the speedups are relatively lower, since the architecture is more compact. We note that the speedups are more significant for batched inference, where the engine has more potential for parallelization, and our setup uses more cores.

B.4 Sparse-Dense Output Comparison

To the best of our knowledge, our results are the first ones to show that *both* a dense and a sparse model can be trained jointly. Although other sparse training techniques such as RigL or Top-KAST can train sparse models faster, none of them offer the additional benefit of an accurate dense model.

Table 14: Comparison with Variational Dropout on ResNet50

Method	Sparsity	Sparse Top-1 Accuracy (%)	Dense Top-1 Accuracy (%)	Sparse-Dense Agreement (%)	Sparse-Dense Cross-entropy
AC/DC	79.83%	76.3 ± 0.1	76.8 ± 0.07	89.8 ± 0.3	0.85 ± 0.005
SparseVD	80%	75.3	75.2	98.6	-
GMP	79.83%	76.4	76.9	86.0	1.03
AC/DC	89.8%	75.0 ± 0.1	76.6 ± 0.09	86.8 ± 1.5	1.02 ± 0.004
SparseVD	90%	73.8	73.6	98.3	-
GMP	89.8%	74.7	76.9	83.5	1.29

One method that does generate sparse-dense model couples is Sparse Variational Dropout (SparseVD) [MAV17]; there, after training a dense model with variational inference, a large proportion of the weights can be pruned in a single step, without affecting the accuracy of the dense model. (We note however that Sparse Variational Dropout doubles the FLOP cost of training, due to the variational parameters.) However, our investigation of the SparseVD models trained by [GEH19] shows that the sparse and dense models agree in over 98% of their predictions as measured on the ImageNet validation set, and are of no better quality than the sparse model - if anything, they are slightly worse. Please see Table 14 for complete results.


In comparison, AC/DC with finetuning produces dense models of validation accuracy that is comparable to that of a dense model trained without any compression, and that therefore do differ from their sparse co-trained counterparts. To understand the relative sizes of these differences, we used GMP pruning as the baseline. In particular, we compared the similarity of a fully trained dense model with GMP trained over 100 epochs. We note AC/DC and GMP show comparable accuracy for both their sparse and dense models in this scenario; however, the total training epochs are substantially lower for producing these models with AC/DC.

We use two metrics to investigate the difference in sparse-dense model pairs: the proportion of validation examples on which the top prediction agreed between the two models, and the average cross-entropy of the predictions across all validation examples. In both metrics, Table 14 shows that model similarity is higher for 80% sparsity than 90%, and higher for AC/DC than GMP training: in particular, sparse/dense cross-entropy is about 20% lower for AC/DC, and the number of top-prediction disagreements is about 25% lower for AC/DC at 90% sparsity, and 37% lower for AC/DC at 80% sparsity.


B.5 Memorization Experiments on CIFAR10

In what follows, we study the similarities between the sparse and dense models learned with AC/DC, on the particular setup of memorizing random labels. Specifically, we select 1000 i.i.d. training samples from the CIFAR10 dataset and randomly change their labels. We train a ResNet20 model using AC/DC, at various target sparsity levels, ranging from 50% to 95%. We use SGD with momentum, weight decay, and initial learning rate 0.1 which is decayed by a factor of 10 every 60 epochs, starting with epoch 65.

Using data augmentation dramatically affects the memorization of randomly-labelled training samples, and thus we differentiate between the two possible cases. Namely, the regular baseline can easily memorize (in the sense of reaching perfect accuracy) the randomly-labelled samples,



(a) Accuracy on the mis-labelled data



(b) Accuracy on the mis-labelled data (w.r.t. the true labels)

Figure 6: Accuracy during training with AC/DC at 90% and 95% target sparsity, for 1000 randomly labelled CIFAR10 images. No data augmentation was applied to the training samples.

when *no* data augmentation is used; in comparison, with data augmentation memorization is more difficult, and the accuracy on randomly-labelled samples for the baseline is just above 60%. In addition to the accuracy on the perturbed samples with respect to their new random labels, we also track the accuracy with respect to the “true” or correct labels. This differentiation offers a better understanding regarding where memorization fails and a glimpse into the robustness properties of neural networks in general, and of AC/DC, in particular.


No data augmentation. As previously mentioned, in this case the baseline model can perfectly memorize the perturbed data, with respect to their random labels. Interestingly, prior to the initial learning rate decay, most ($\geq 70\%$) perturbed samples are still correctly classified with respect to their “true” labels, and memorization happens very quickly after the learning rate is decreased. In the case of AC/DC with low target sparsity (50% and 75%), memorization has a very similar behavior to the dense baseline. However, for higher sparsity levels (90% and 95%) we can see a clear difference between the sparse and dense models. Namely, during each compression phase most perturbed samples are correctly classified with respect to their true labels, whereas in decompression phases their random labels are memorized. This phenomenon is illustrated in Figure 6.

Data augmentation. In this case, memorization of the perturbed samples is more difficult, and it happens later on during training, usually after the second learning rate decrease for the baseline model. Interestingly, in the case of AC/DC we can see (Figure 7) a clear inverse relationship between the amount of memorization and the target sparsity. Although low sparsity enables more memorization, most perturbed samples are still correctly classified with respect to their true labels. For higher sparsity levels (90% and 95%), most perturbed samples are correctly classified with respect to their true labels (almost 90%) and very few are memorized. Furthermore, the dense model resulted from AC/DC training is more robust than the original baseline, as it still learns the correct labels of the perturbed samples, despite being presented with random ones.


C Computational Details

C.1 FLOPs Computation

When computing FLOPs, we take into account the number of zero-valued weights for linear and convolutional layers. We also include in our calculations FLOPs for Batch Normalization, Pooling and ReLU layers, although they have a negligible impact towards the final number. To compute the FLOPs required for a backward pass over a sample, we use the same convention as RigL [EGM⁺20];



(a) Accuracy on the mis-labelled data



(b) Accuracy on the mis-labelled data (w.r.t. the true labels)

Figure 7: Accuracy during training with AC/DC at 50%, 75%, 90% and 95% target sparsity, for 1000 randomly labelled CIFAR10 images. Here, all samples were trained using data augmentation.

namely, if F denotes the inference FLOPs per sample, the number of backward FLOPs is estimated as $B = 2 \cdot F$, as we need F FLOPs to backpropagate the error, and additional F to compute the gradients w.r.t. the weights.

For compression and decompression phases C and D , we consider F_C and F_D the compression and decompression inference FLOPs per sample, respectively. We use F to denote the inference FLOPs per sample for the baseline network. During each compression phase, the training FLOPs per sample can be estimated as $3 \cdot F_C$. For decompression phases, we noticed that a small fraction of weights remain zero, and therefore $F_D < F$. When doing a backward pass we have additional F_D from back-propagating the error, and F extra FLOPs for the gradients with respect to all parameters. Therefore, we estimate the training FLOPs per sample during a decompression phase as $2 \cdot F_D + F$. To obtain the final number of FLOPs, we compute the inference FLOPs on a random input sample, estimate the backward FLOPs, compute the estimated training FLOPs over all training epochs as described above, and scale by the number of training samples.

C.2 Choice of Hyper-parameters

Length of compression/decompression phases. AC/DC alternates between compression and decompression phases to co-train sparse and dense models. It is important to note, however, that the length of these phases, together with the warm-up and fine-tuning phases, could have a significant impact on the quality of the resulting models. Before settling on the sparsity pattern we used for all our ImageNet experiments (see Figure 2a and Figure 5a), we experimented with different lengths for the sparse/dense phases, but found that ultimately the pattern used in the paper had the best trade-off between training FLOPs and validation accuracy.

Notably, we experimented on ResNet50, 90% sparsity, with increasing the training epochs to 130 and with different lengths for the compression/decompression phases. For example, we found that alternating between sparse/dense training every 10 epochs yielded slightly better results after 130 epochs: 75.34% for the sparse model, 76.87% for the fine-tuned dense model; however, this also had higher training FLOPs requirements ($0.7\times$ for the sparse model and $0.9\times$ including the fine-tuned dense). We additionally experimented with longer dense phases (10 epochs), compared to sparse phases (5 epochs); this also resulted in more accurate models: 75.45% accuracy for the sparse model and 76.78% – for the fine-tuned dense model. However, the training FLOPs were substantially higher: $0.85\times$ for the sparse model and $1.15\times$ for the fine-tuned dense.

Due to computational limitations, and to ensure a fair comparison with the dense baseline and other pruning methods, we decided on using a fixed number of 100 training epochs (the same used for the dense baseline). In this setup, we did not experiment with other lengths for the compression/decompression phases other than the ones used in Figure 2a and Figure 5a, but noticed that having a longer final decompression phase had a positive impact on the fine-tuned dense model. For instance, when following a sparsity schedule as in Figure 1, the sparse model at 90% sparsity had a very similar performance to the reported results (75.18% accuracy, from one seed), while the fine-tuned dense model was significantly below the dense baseline (76.05% validation accuracy). In general, we advise simplicity in designing the sparsity schedule, by keeping the same length for the sparse/dense phases. We believe having a short warm-up period and a longer fine-tuning phase are both beneficial for the sparse model; in our experiments, we only used warm-up phases of 10 epochs, but believe that shorter phases are worth exploring as well. Furthermore, the mask difference between consecutive compression phases is an important guide for choosing the sparsity schedule: as it was previously discussed, having a non-trivial difference between the masks typically results in better sparse models. Illustrations of the pruning masks during training on ImageNet are presented in Figure 2b and Figure 5b.

When choosing the sparsity schedule for the language models experiments on Transformers-XL, we followed the same principles as for ImageNet. In fact, the sparsity schedule is very similar to the one used for ImageNet, scaled by the number of training epochs (48 epochs or 100,000 steps for Transformers).

In the case of CIFAR100, we used for AC/DC the same number of 200 training epochs as for the dense baseline. We experimented with sparse/dense phases of lengths 10 or 20, and found that generally switching every 20 epochs between sparse and dense training yielded the best results.

Training Hyper-parameters for ImageNet. We used the same hyper-parameters for all our ImageNet experiments, on both ResNet50 and MobileNetV1. Namely, we trained using SGD with momentum and batch size 256. We used a cosine learning rate scheduler, after an initial warm-up phase of 5 epochs, when the learning rate was linearly increased to 0.256. The momentum value was 0.875 and weight decay was 0.00003051757813. These hyper-parameters have the standard values used in the implementation of STR [KRS⁺20]. Furthermore, to improve efficiency, we train and evaluate the models using mixed precision (FP16); the difference in accuracy versus FP32 is negligible (<0.05%).

Training Hyper-parameters for Transformer-XL. For our Transformer-XL experiments, we integrated into our code-base the implementation provided by NVIDIA ¹, which also follows closely the original implementation in [DYY⁺19]. We used the same hyper-parameters for training the large Transformer-XL model with 18 layers on WikiText-103, including the Lamb optimizer [YLR⁺19] with cosine learning rate scheduler.

References

- [AHJ⁺18] Dan Alistarh, Torsten Hoefler, Mikael Johansson, Sarit Khirirat, Nikola Konstantinov, and Cédric Renggli. The convergence of sparsified gradient methods. In *Advances in Neural Information Processing Systems*, 2018.
- [AS20] Kyriakos Axiotis and Maxim Sviridenko. Sparse convex optimization via adaptively regularized

¹<https://github.com/NVIDIA/DeepLearningExamples/tree/master/PyTorch/LanguageModeling/Transformer-XL>

- hard thresholding. In *International Conference on Machine Learning*, pages 452–462. PMLR, 2020.
- [AZLS19] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pages 242–252. PMLR, 2019.
- [BD08] Thomas Blumensath and Mike E Davies. Iterative thresholding for sparse approximations. *Journal of Fourier analysis and Applications*, 14(5-6):629–654, 2008.
- [Bie20] Lukas Biewald. Experiment tracking with weights and biases, 2020. Software available from wandb.com.
- [BKML17] Guillaume Bellec, David Kappel, Wolfgang Maass, and Robert Legenstein. Deep rewiring: Training very sparse deep networks. *arXiv preprint arXiv:1711.05136*, 2017.
- [CHS⁺16] Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1. *arXiv preprint arXiv:1602.02830*, 2016.
- [CMJ⁺18] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, et al. TVM: An automated end-to-end optimizing compiler for deep learning. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, pages 578–594, 2018.
- [CT06] Emmanuel J Candes and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE transactions on information theory*, 52(12):5406–5425, 2006.
- [DCP17] Xin Dong, Shangyu Chen, and Sinno Jialin Pan. Learning to prune deep neural networks via layer-wise optimal brain surgeon, 2017.
- [DDJ⁺20] Robert David, Jared Duke, Advait Jain, Vijay Janapa Reddi, Nat Jeffries, Jian Li, Nick Kreeger, Ian Nappier, Meghna Natraj, Shlomi Regev, et al. Tensorflow lite micro: Embedded machine learning on tinymml systems. *arXiv preprint arXiv:2010.08678*, 2020.
- [Dee21] DeepSparse. Neuralmagic deepsparse inference engine, 2021.
- [DYY⁺19] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019.
- [DZ19] Tim Dettmers and Luke Zettlemoyer. Sparse networks from scratch: Faster training without losing performance. *arXiv preprint arXiv:1907.04840*, 2019.
- [EDGS20] Erich Elsen, Marat Dukhan, Trevor Gale, and Karen Simonyan. Fast sparse convnets. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14629–14638, 2020.
- [EGM⁺20] 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.
- [FC18] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2018.
- [FDRC19] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M Roy, and Michael Carbin. Stabilizing the lottery ticket hypothesis. *arXiv preprint arXiv:1903.01611*, 2019.
- [Fou11] Simon Foucart. Hard thresholding pursuit: an algorithm for compressive sensing. *SIAM Journal on Numerical Analysis*, 49(6):2543–2563, 2011.

- [Fou12] Simon Foucart. Sparse recovery algorithms: sufficient conditions in terms of restricted isometry constants. In *Approximation Theory XIII: San Antonio 2010*, pages 65–77. Springer, 2012.
- [GEH19] Trevor Gale, Erich Elsen, and Sara Hooker. The state of sparsity in deep neural networks. *arXiv preprint arXiv:1902.09574*, 2019.
- [GKD⁺21] Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W Mahoney, and Kurt Keutzer. A survey of quantization methods for efficient neural network inference. *arXiv preprint arXiv:2103.13630*, 2021.
- [Gra21] Graphcore. Graphcore poplar sdk 2.0, 2021.
- [HABN⁺21] Torsten Hoefer, Dan Alistarh, Tal Ben-Nun, Nikoli Dryden, and Alexandra Peste. Sparsity in deep learning: Pruning and growth for efficient inference and training in neural networks. *arXiv preprint arXiv:2102.00554*, 2021.
- [Hag94] Masafumi Hagiwara. A simple and effective method for removal of hidden units and weights. *Neurocomputing*, 6(2):207 – 218, 1994. Backpropagation, Part IV.
- [HPN⁺16] Song Han, Jeff Pool, Sharan Narang, Huizi Mao, Enhao Gong, Shijian Tang, Erich Elsen, Peter Vajda, Manohar Paluri, John Tran, et al. Dsd: Dense-sparse-dense training for deep neural networks. *arXiv preprint arXiv:1607.04381*, 2016.
- [HPTD15] Song Han, Jeff Pool, John Tran, and William J Dally. Learning both weights and connections for efficient neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 1*, pages 1135–1143, 2015.
- [HSW93] Babak Hassibi, David G Stork, and Gregory J Wolff. Optimal brain surgeon and general network pruning. In *IEEE international conference on neural networks*, pages 293–299. IEEE, 1993.
- [HZC⁺17] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [HZRS16] 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.
- [JPR⁺20] 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.
- [JTK14] Prateek Jain, Ambuj Tewari, and Purushottam Kar. On iterative hard thresholding methods for high-dimensional M-estimation. In *Proceedings of the 27th International Conference on Neural Information Processing Systems-Volume 1*, pages 685–693, 2014.
- [JYFY16] Xiaojie Jin, Xiaotong Yuan, Jiashi Feng, and Shuicheng Yan. Training skinny deep neural networks with iterative hard thresholding methods. *arXiv preprint arXiv:1607.05423*, 2016.
- [KB14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [KH⁺09] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [KNS16] Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 795–811. Springer, 2016.

- [KRS⁺20] Aditya Kusupati, Vivek Ramanujan, Raghav Somani, Mitchell Wortsman, Prateek Jain, Sham Kakade, and Ali Farhadi. Soft threshold weight reparameterization for learnable sparsity. In *International Conference on Machine Learning*, pages 5544–5555. PMLR, 2020.
- [LAT18] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip HS Torr. Snip: Single-shot network pruning based on connection sensitivity. *arXiv preprint arXiv:1810.02340*, 2018.
- [LDS90] Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. In *Advances in Neural Information Processing Systems*, pages 598–605, 1990.
- [LSB⁺19] Tao Lin, Sebastian U Stich, Luis Barba, Daniil Dmitriev, and Martin Jaggi. Dynamic model pruning with feedback. In *International Conference on Learning Representations*, 2019.
- [LZB20] Chaoyue Liu, Libin Zhu, and Mikhail Belkin. Toward a theory of optimization for over-parameterized systems of non-linear equations: the lessons of deep learning. *arXiv preprint arXiv:2003.00307*, 2020.
- [MAV17] Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov. Variational dropout sparsifies deep neural networks. In *International Conference on Machine Learning*, pages 2498–2507. PMLR, 2017.
- [MLP⁺21] Asit Mishra, Jorge Albericio Latorre, Jeff Pool, Darko Stosic, Dusan Stosic, Ganesh Venkatesh, Chong Yu, and Paulius Micikevicius. Accelerating sparse deep neural networks. *arXiv preprint arXiv:2104.08378*, 2021.
- [MMS⁺18] 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):1–12, 2018.
- [MXBS16] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016.
- [PGM⁺19] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*. 2019.
- [Qia99] Ning Qian. On the momentum term in gradient descent learning algorithms. *Neural networks*, 12(1):145–151, 1999.
- [RDS⁺15] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [SA20] Sidak Pal Singh and Dan Alistarh. Woodfisher: Efficient second-order approximation for neural network compression. *Advances in Neural Information Processing Systems*, 33, 2020.
- [SCCS19] Shaohuai Shi, Xiaowen Chu, Ka Chun Cheung, and Simon See. Understanding top-k sparsification in distributed deep learning. *arXiv preprint arXiv:1911.08772*, 2019.
- [SGM20] Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for modern deep learning research. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13693–13696, 2020.
- [SM20] Alexander Shevchenko and Marco Mondelli. Landscape connectivity and dropout stability of sgd solutions for over-parameterized neural networks. In *International Conference on Machine Learning*, pages 8773–8784. PMLR, 2020.

- [TKYG20] 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, 2020.
- [Van17] Han Vanholder. Efficient inference with TensorRT. NVIDIA GTC On-Demand. Slides available at <https://on-demand-gtc.gputechconf.com/gtcnew/sessionview.php?sessionName=23425-efficient+inference+with+tensorrt>, 2017.
- [VSP⁺17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6000–6010, 2017.
- [YLR⁺19] Yang You, Jing Li, Sashank Reddi, Jonathan Hseu, Sanjiv Kumar, Srinadh Bhojanapalli, Xiaodan Song, James Demmel, Kurt Keutzer, and Cho-Jui Hsieh. Large batch optimization for deep learning: Training bert in 76 minutes. In *International Conference on Learning Representations*, 2019.
- [YLZ14] Xiaotong Yuan, Ping Li, and Tong Zhang. Gradient hard thresholding pursuit for sparsity-constrained optimization. In *International Conference on Machine Learning*, pages 127–135. PMLR, 2014.
- [ZBH⁺16] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.
- [ZG17] 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.
- [ZK16] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *British Machine Vision Conference 2016*. British Machine Vision Association, 2016.
- [ZMZ⁺21] 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. *arXiv preprint arXiv:2102.04010*, 2021.