

Structured Pruning of Large Language Models

Ziheng Wang*
ASAPP, Inc.

zihengw@stanford.edu

Jeremy Wohlwend*
ASAPP, Inc.

jwohlwend@csail.mit.edu

Tao Lei*
ASAPP, Inc.

tao@asapp.com

Abstract

Large language models have recently achieved state of the art performance across a wide variety of natural language tasks. Meanwhile, the size of these models and their latency have significantly increased, which makes their usage costly, and raises an interesting question: do language models need to be large? We study this question through the lens of model compression. We present a generic, structured pruning approach by parameterizing each weight matrix using its low-rank factorization, and adaptively removing rank-1 components during training. On language modeling tasks, our structured approach outperforms other unstructured and block-structured pruning baselines at various compression levels, while achieving significant speedups during both training and inference. We also demonstrate that our method can be applied to pruning adaptive word embeddings in large language models, and to pruning the BERT model on several downstream fine-tuning classification benchmarks.¹

1 Introduction

Recent advances in language modeling have led to remarkable improvements on a variety of natural language tasks (Dai and Le, 2015; Peters et al., 2018; Radford et al., 2018; Devlin et al., 2019; Liu et al., 2019; Dai et al., 2019; Zhang et al., 2019). These models, however, have grown increasingly large, rendering them slow and expensive for real-world applications. Through the use of model compression, we aim to reduce this overhead, and to better understand the role of model capacity in large language models.

A common approach to model compression is known as weight pruning (Zhu and Gupta, 2017; Han et al., 2015a; See et al., 2016). Model weights

are progressively removed, resulting in sparse matrices across the network. Earlier work focuses mostly on unstructured pruning, where weights are pruned individually (Narang et al., 2017a; Zhu and Gupta, 2017). While this method is effective, it results in unstructured sparse matrices that are difficult to support on common hardware (Han et al., 2016), making it challenging to obtain training and inference speedups despite a significant reduction in model size.

On the other hand, *structured pruning* imposes structured sparse patterns by removing groups of consecutive parameters, such as rows, columns or $k \times k$ sub-blocks of the weight matrix (Narang et al., 2017b; Wen et al., 2018; Cao et al., 2019). These methods lead to significant speedup, but tend to give lower performance than unstructured pruning given the same parameter budget (Yao et al., 2019). Another caveat is that some of these methods require special linear algebra implementations (Gray et al., 2017; Yao et al., 2019) or hardware (Cao et al., 2019) in order to accelerate matrix multiplication, therefore limiting their application to a broad set of existing models.

We propose a generic, improved structured pruning approach based on adaptive low-rank factorization. As an alternative to unstructured sparse and block sparse representations, low-rank factorization retains the full dense structure of weight matrices, eliminating the need for special linear algebra primitives and hardware for computation speedup. Compared to row (and column) based pruning, low-rank factorization better preserves the linear transformation of the un-compressed matrices. During training, our method adaptively learns which low-rank components to remove in order to achieve a strong performance-compression trade-off. We show that a simple magnitude based pruning strategy is sufficient to accomplish strong results. In addition, we further increase performance via an improved l_0 regularization (Louizos et al., 2018)

*Denotes equal contribution.

¹Our code is publicly available at <https://github.com/asappresearch/flop>.

technique which uses an augmented Lagrangian method to directly control the final compression level of the model. Our method, which we refer to as FLOP (Factorized Low-rank Pruning) applies to any matrix multiplication.

Pruning large language models introduces unique challenges with the handling of large input and output layers. Although our method is generic, it is particularly well suited to this task. In particular, we show that FLOP can dynamically learn the embedding dimensions of different word clusters, effectively extending the idea of adaptive embeddings and softmax (Grave et al., 2017; Baevski and Auli, 2019). Since these embedding layers take a significant amount of parameters in the language models, learning flexible dimensions instead of specifying them manually results in a more optimal trade-off between parameter reduction and performance.

We evaluate our method on common language modeling and language understanding tasks including the Wiki-103, Enwiki8 and GLUE benchmarks, and by testing our method on both recurrent networks and Transformer (Vaswani et al., 2017). Our results demonstrate that factorization based pruning significantly outperforms block-structured pruning and even surpasses unstructured pruning, while using our improved l_0 regularization further improves the performance in most cases. When pruning a large word-level language model with adaptive embeddings for example, our method achieves 50% compression while losing only 0.8 perplexity. Moreover, our method is able to achieve over 2x speed-up during both training and inference with no additional hardware or software requirements. Our method will be released as a Pytorch (Paszke et al., 2017) library.

2 Related Work

The development of model compression techniques can be categorized into three areas of research: weight pruning (Han et al., 2015b; Zhu and Gupta, 2017), knowledge distillation (Ba and Caruana, 2014; Hinton et al., 2015; Kim and Rush, 2016), and quantization (Gong et al., 2014; Zhu et al., 2017; Shen et al., 2019).

Recent efforts have successfully applied compression on various architectures and NLP applications, such as pruning multi-head attentions for machine translation (Voita et al., 2019), learning adaptive embeddings and softmax layers for lan-

guage models (Grave et al., 2017; Baevski and Auli, 2019; Li et al., 2018; Variani et al., 2019), and compressing BERT models via distillation (Chia et al., 2019; Jiao et al., 2019; Sanh et al., 2019; Sun et al., 2019; Tsai et al., 2019; Turc et al., 2019). Only one of the compression techniques such as distillation has been used in these works for simplicity. However, these techniques can be combined to achieve greater compression (Han et al., 2015a; Shangguan et al., 2019). Our pruning method is compatible with quantization and distillation, as it can be applied to compress any matrix multiplication in a network.

Previous work has considered different weight pruning approaches such as unstructured pruning based on magnitude (Narang et al., 2017a; Frankle and Carbin, 2019), dropout (Gale et al., 2019; Fan et al., 2020; Molchanov et al., 2017), and structured pruning (Wen et al., 2018; Louizos et al., 2017). Model weights are often removed via thresholding and l_1 regularization during the pruning process (Narang et al., 2017b; Liu et al., 2018). Our method differs from previous work by using low-rank parameterization for compression. Furthermore, we extend l_0 regularization using an augmented Lagrangian optimization method to control the final model size.

3 Background

We formalize the task of model pruning as an end-to-end learning problem with l_0 regularization, following the prior work of Louizos et al. (2018).

Consider a given neural network model $f(\cdot; \theta)$ parameterized by $\theta = \{\theta_j\}_{j=1}^n$, where each θ_j represents an individual parameter weight or a block of weights (e.g. a column of a weight matrix) and n denotes the number of blocks. A pruning strategy of the model can be parameterized by introducing additional binary variables $\mathbf{z} = \{z_j\}_{j=1}^n$ such that $z_j \in \{0, 1\}$ and

$$\tilde{\theta} = \theta \odot \mathbf{z} \quad \forall j \quad \tilde{\theta}_j = \theta_j z_j.$$

Here $\tilde{\theta} = \{\tilde{\theta}_j\}$ denotes the set of model parameters after pruning and its l_0 norm, $\|\tilde{\theta}\|_0 = \sum_{j=1}^n z_j$, measures the effective size of the pruned model.

The choice of binary variables \mathbf{z} can be regulated by some prior distribution and optimized given the training data. That is, let $q_j(z)$ be the density function of the learnable prior of z_j . The optimization objective during training can be formulated as min-

imizing the expected training loss

$$\mathbb{E}_{\mathbf{z}} \left[\frac{1}{D} \sum_{i=1}^D \mathcal{L}(\mathbf{x}_i, \mathbf{y}_i; \tilde{\boldsymbol{\theta}}) + \lambda \|\tilde{\boldsymbol{\theta}}\|_0 \right], \quad (1)$$

where $\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^D$ are training examples, \mathcal{L} is the training loss function and $\lambda > 0$ is a constant hyperparameter for l_0 norm regularization encouraging the model to be sparse. Note that in practice optimizing this objective is intractable due to the discrete nature of z_j and an exponential number of 2^n choices.

The key to the method of Louizos et al. (2018), called the re-parameterization trick, enables \mathbf{z} to be differentiable and jointly trained with the model parameter $\boldsymbol{\theta}$. Specifically, the random variables \mathbf{z} are relaxed as continuous variables distributed within the interval $[0, 1]$. In addition, instead of learning the probability density function $q_j(z)$, the re-parameterization trick proposes to learn the inverse of the cumulative density function (CDF). Note that if $G(\cdot)$ is the inverse of CDF for a variable z , then z can be easily sampled by first sampling $u \sim U(0, 1)$ and computing $z = G(u)$. Assuming the inverse CDF function is parameterized by some learnable parameters $\boldsymbol{\alpha} = \{\alpha_j\}_{j=1}^n$ and the function $G(\cdot; \boldsymbol{\alpha})$ is differentiable, we obtain an overall end-to-end learning objective,

$$\min_{\boldsymbol{\theta}, \boldsymbol{\alpha}} \mathbb{E}_{\mathbf{u} \sim U(0,1)} \left[\frac{1}{D} \sum_{i=1}^D \mathcal{L}(\mathbf{x}_i, \mathbf{y}_i; \tilde{\boldsymbol{\theta}}) + \lambda \|\tilde{\boldsymbol{\theta}}\|_0 \right], \quad (2)$$

$$z_j = G(u_j; \alpha_j), \quad \forall j = 1 \cdots n$$

where $\mathbf{u} = \{u_1, \dots, u_n\}$ denotes the iid samples from the uniform distribution. Since \mathbf{z} is now the output of the parameterized function $G(\cdot; \boldsymbol{\alpha})$ and is used as an intermediate representation for the neural network (with $\tilde{\boldsymbol{\theta}} = \boldsymbol{\theta} \odot \mathbf{z}$), gradient based optimization methods can perform gradient updates for $\boldsymbol{\theta}$ and $\boldsymbol{\alpha}$.

Following previous work, we choose the Hard Concrete distribution for the random variables $\mathbf{z} = \{z_j\}$. The inverse of CDF $G(\cdot; \boldsymbol{\alpha})$ of this distribution is defined as follows

$$\begin{aligned} \mathbf{u} &\sim U(0, 1) \\ \mathbf{s} &= \text{sigmoid}(\log \mathbf{u} - \log(1 - \mathbf{u}) + \boldsymbol{\alpha}) \\ \bar{\mathbf{s}} &= \mathbf{s} \times (r - l) + l \\ \mathbf{z} &= \min(1, \max(0, \bar{\mathbf{s}})) \end{aligned}$$

where $l < 0$ and $r > 1$ are two constants used to ‘stretch’ the sigmoid outputs \mathbf{s} into the interval

(l, r) , and the final outputs \mathbf{z} are rectified into $[0, 1]$. The stretch-and-rectify process has the effect of assigning a significant portion of probability mass on the integer values $\{0, 1\}$, which makes it a good relaxation of the binary (Bernoulli) distribution. During training, we sample \mathbf{u} and compute \mathbf{z} and the loss $\mathcal{L}(\cdot)$ for each training batch. The expected l_0 norm regularization can be separately computed via a closed form

$$\begin{aligned} \mathbb{E} \left[\|\tilde{\boldsymbol{\theta}}\|_0 \right] &= \sum_{j=1}^n \mathbb{E} [z_j > 0] \\ &= \sum_{j=1}^n \text{sigmoid} \left(\alpha_j - \log \frac{-l}{r} \right) \end{aligned} \quad (3)$$

which is differentiable as well.

4 Method

In this section, we introduce FLOP, an improved structured pruning method. FLOP proposes a different parameterization of the weight matrices using low-rank factorization. In addition, we introduce a revised optimization objective that allows for an explicit control of the compression size.

4.1 Structured Pruning using Factorization

In weight pruning, a key choice is how we define parameter blocks $\theta_1, \dots, \theta_n$ to achieve the most effective pruning results. One obvious method is to prune each individual parameter weight, which often retains strong performance but poses challenges to achieve a computation speedup given unstructured sparse matrices.

Structured pruning chooses to remove groups of consecutive parameters as a remedy. For example, consider a fully connected layer which performs a multiplication $\mathbf{W}\mathbf{x}$ for an input feature $\mathbf{x} \in \mathbb{R}^d$ and weight matrix $\mathbf{W} \in \mathbb{R}^{d' \times d}$. One popular method, sometimes referred to as neuron or input feature pruning, consists of adding the sparsity variables as a sparse diagonal matrix $\mathbf{G} = \text{diag}(z_1, \dots, z_d)$ to the multiplication, i.e., $\mathbf{W}\mathbf{G}\mathbf{x}$. This effectively removes the subset of the columns in \mathbf{W} with $z_k = 0$, where k is the column index. In practice, this method produces significant speedups at both training and inference time (by selecting a small subset of columns and performing matrix multiplications given much smaller matrices). However, it is reported to achieve lower performance compared to unstructured pruning (Yao et al., 2019) due to more restrictive sparse patterns.

We propose to use low-rank factorization as a less restrictive, yet powerful representation and obtain parameter reduction by pruning rank-1 components. That is, we reparameterize and factorize the matrix \mathbf{W} into the product of two smaller matrices $\mathbf{W} = \mathbf{P}\mathbf{Q}$, where $\mathbf{P} \in \mathbb{R}^{d' \times r}$, $\mathbf{Q} \in \mathbb{R}^{r \times d}$ and $r \leq \min\{d, d'\}$ is the number of columns of \mathbf{P} (equivalently the number of rows of \mathbf{Q}). Let \mathbf{p}_k and \mathbf{q}_k be the k -th column of \mathbf{P} and k -th row of \mathbf{Q} respectively. Since \mathbf{W} is now the sum of r rank-1 components $\mathbf{p}_k \mathbf{q}_k$, we can achieve structured pruning by introducing a pruning variable z_k for each component

$$\mathbf{W} = \mathbf{P}\mathbf{G}\mathbf{Q} = \sum_{k=1}^r z_k \times (\mathbf{p}_k \times \mathbf{q}_k)$$

where $\mathbf{G} = \text{diag}(z_1, \dots, z_r)$ is again a diagonal matrix of pruning variables. Intuitively, learning the factorization has the potential of keeping the most effective rank-1 components, and thereby better preserve the model performance.¹

After training, only columns and rows corresponding to non-zero diagonal values need to be stored, resulting in much smaller (but still dense) matrices. The nonzero values of \mathbf{G} can be absorbed into either \mathbf{P} or \mathbf{Q} . The computation boils down to simple matrix multiplications at inference time, maximizing efficiency on common hardware. Unlike unstructured pruning, we need not store the indices of the sparse weights, resulting in greater memory savings.

4.2 Pruning Adaptive Embedding and Softmax Layer

The input embedding and softmax output layer can take the vast majority of parameters in a language model when the vocabulary size is large. Previous work have considered various techniques that are specifically tailored to compress the embedding and softmax layer. For instance, the adaptive embedding and softmax methods of Grave et al. (2017); Baevski and Auli (2019) have been shown to achieve impressive results in preserving perplexity while significantly reducing the total number of embedding parameters.

We describe how FLOP fits naturally with these adaptive methods, giving them more potential. The core idea behind the adaptive methods is to apply

¹It is also easy to see that input feature pruning is a special case of low-rank pruning: By fixing $\mathbf{P} = \mathbf{W}$ and $\mathbf{Q} = \mathbf{I}$, $\mathbf{P}\mathbf{G}\mathbf{Q} = \mathbf{W}\mathbf{G}\mathbf{I} = \mathbf{W}\mathbf{G}$.

different embedding dimensions and projections to different word clusters. Consider the recent method of Baevski and Auli (2019) without loss of generality. Let $i \in \{1, \dots, C\}$ denotes the indice of the i -th word cluster (sorted based on word frequency). Two parameter matrices $\mathbf{E}_i \in \mathbb{R}^{n_i \times d_i}$ and $\mathbf{O}_i \in \mathbb{R}^{d_i \times d}$ are introduced for the i -th cluster, where n_i is the number of words in the cluster, d is the original embedding dimension and d_i is the reduced word dimension for this cluster. In other words, each word embedding in this cluster has dimension d_i but are projected back into dimension d using a projection \mathbf{O}_i (and vice versa). This is in indeed a low-rank factorization

$$\tilde{\mathbf{E}}_i = \mathbf{E}_i \mathbf{O}_i \in \mathbb{R}^{n_i \times d}$$

for an underlying embedding matrix $\tilde{\mathbf{E}}_i$. While the reduced dimensions $\{d_i\}_{i=1}^C$ usually have to be manually specified, our method automatically learns separate diagonal pruning mask \mathbf{G}_i for each cluster, i.e. $\tilde{\mathbf{E}}_i = \mathbf{E}_i \mathbf{G}_i \mathbf{O}_i$. During training and pruning, it *adaptively* learns to adjust the parameter budget of each word cluster based on what is needed to achieve good performance. Unsurprisingly, our method prunes most of the dimensions for rare words, which is consistent with the empirical choice made in prior work.

4.3 Augmented Lagrangian Method

Our method can be implemented with a magnitude based pruning strategy, or directly trained with the training objective (2) which uses an l_0 regularization $\lambda \|\tilde{\theta}\|_0$ to promote weight pruning. One limitation of this regularization however is the lack of effective control on the size of the pruned model. For instance, we observe that training with the same λ could converge to very different model sizes when using slightly different learning rates or pruning schedules. This can be problematic because a desired model size or parameter budget is often needed in many real-world applications.

We make use of an Augmented Lagrangian method to overcome this training limitation. Lagrangian relaxation methods have been explored in many NLP problems (Bastings et al., 2019; Martins et al., 2011; Flanigan et al., 2014; Rush et al., 2010). We use the following Lagrangian variant for our task – Let t be the target model size and $s(\alpha)$ be the expected model size determined by the Hard Concrete parameter α . Note $s(\alpha)$ can be computed based on Eq (3) by multiplying $\mathbb{E}[z_j > 0]$ with the

size of the j -th parameter block. Our Augmented Lagrangian method imposes an equality constraint $s(\alpha) = t$ by introducing a violation penalty,

$$g(\lambda, \alpha) = \lambda_1 \cdot (s(\alpha) - t) + \lambda_2 \cdot (s(\alpha) - t)^2$$

where $\lambda_1, \lambda_2 \in \mathbb{R}$ are two Lagrangian multipliers that will be jointly updated during training. The overall training optimization is an adversarial game,

$$\max_{\lambda_1, \lambda_2} \min_{\theta, \alpha} \mathbb{E}_{\mathbf{u}} \left[\frac{1}{D} \sum_{i=1}^D \mathcal{L}(\mathbf{x}_i, \mathbf{y}_i; \tilde{\theta}) \right] + g(\lambda, \alpha).$$

The updates of λ_1 and λ_2 would always increase the training loss unless the equality constraint is met, which gives us the desired model size.

We gradually increase the target size t at a linear rate during the process of pruning training. That is, given the desired size t_{\max} , we set the sparsity at k -th pruning iteration as

$$t_k = \min\left(1, \frac{k}{m}\right) \cdot t_{\max}$$

where m is a hyperparameter specifying the number of annealing steps.

We perform joint gradient updates for the model parameters θ, α as well as the Lagrangian multipliers λ_1, λ_2 . For each training batch, we sample the pruning mask $\mathbf{z} = \{z_1, \dots, z_n\}$ and share it across the training examples within the batch. Since the pruning mask is shared, we can select parameters that are only active for the current batch and compute smaller matrix multiplications in forward and backward passes. This results in training speedup when \mathbf{z} becomes sparse.

4.4 Inference

During training, the prune mask is a random variable drawn from the Hard Concrete distribution. At inference time, however, we must use a deterministic, fixed mask \mathbf{z} for each weight matrix to obtain the compressed factorization matrices \mathbf{P} and \mathbf{Q} (by keeping i -th low-rank component if $z_i > 0$). We do so by computing the expected value of each z_i in \mathbf{z} using Eq.(3) described in Section 3, and then keeping the top values of $\{z_1, \dots, z_n\}$ and clipping the rest to zero, as to match the l_0 norm (i.e. the compression level).

5 Experimental Setup

Tasks We evaluate the performance of our method on language modeling and BERT fine-

tuning. Specifically, we consider the following task setup.

1. Recurrent word-level language models on the Wiki-103 dataset. We adopt SRU (Lei et al., 2018) as the recurrent architecture and tied adaptive embedding and softmax layers (Baeovski and Auli, 2019). Our base model consists of 12 recurrent layers, 100M parameters in total. About 50% of the parameters are used for the adaptive layers.
2. Recurrent character-level language models on the Enwik8 dataset. We use the same SRU architecture. The base model uses 6 recurrent layers and 35M parameters in total.
3. Transformer-XL model on the Enwik8 dataset. We use the 12-layer base model from Dai et al. (2019) containing 41M parameters. We introduce pruning for the matrices in the self-attention layers as well as those in the feed-forward layers. For factorization based pruning, we choose the starting rank r for each matrix such that the number of parameters remain the same as the unfactorized model².
4. BERT fine-tuning on several classification benchmarks benchmark (Socher et al., 2013; Dolan and Brockett, 2005; Cer et al., 2017; Wang et al., 2019). In this experiment, we use the pre-trained RoBERTa base model by Liu et al. (2019).

We extend the implementation of Transformer, SRU and the adaptive embedding / softmax layers to support factorization based pruning (and other baselines).

Baselines We compare with the following unstructured, structured and/or factorization based pruning baselines.

- **FAC** which trains low-rank factorized models from scratch by reducing all dimensions with the same ratio to get the desired compression.
- **NP- l_0** (Louizos et al., 2018) which adopts l_0 regularization and performs neuron pruning (i.e. removing input features and columns of weight matrices). No factorization is used for this baseline. We add the Augmented Lagrangian optimization similar to FLOP to achieve the exact desired compression.

²In effect, we set $r = d_1 d_2 / (d_1 + d_2)$, where d_1, d_2 are the dimensions of the original weight matrix.

- **AGP** (Zhu and Gupta, 2017) which gradually prunes individual parameters based on the weight magnitude. AGP is one of the state-of-the-art unstructured pruning methods. We use the implementation provided in the Nervana Distiller library (Zmora et al., 2019).
- **FLOP-AGP** is a variant of our full method that prunes low-rank components, but uses magnitude-based gradual pruning on the diagonal mask \mathbf{G} instead. We also tune l_1 regularization on the masks to encourage sparsity, similar to Narang et al. (2017b).

These baselines serve as competitive pruning alternatives, and also provide data points for us to isolate the effectiveness of sub-components of our method, such as low-rank factorization and l_0 pruning. All methods use the same training configurations such as learning rate and dropout. We tune hyper-parameters related to pruning such as compression scheduling and the learning rate of Lagrangian variables for each method. More training and implementation details are provided in the appendix.

6 Results

Word-level Language Model Table 1 presents the results of FLOP as well as the baseline methods. The SRU base model (unpruned) achieves a test perplexity of 24.5, being a strong starting point and competitive with top-performing models such as Transformer (Dai et al., 2019).

The pruning results conform to our expectations that pruning a large model is consistently better than training a small model from scratch, and using low-rank based pruning yields better performance than removing matrix columns and input features. FLOP exceeds the performance of FAC, NP- l_0 and AGP baselines at all compression levels tested. The performance of FLOP-AGP, especially in comparison with its unstructured counterpart AGP, highlights the effectiveness of factorization based pruning. Moreover, we achieve a test perplexity (PPL) of 25.3 with FLOP- l_0 method, a loss of 0.8 perplexity score, while removing 50% of the model parameters. This result is impressive since our base model adopts the adaptive word embedding and softmax layers, which already reduce the model size significantly.

Figure 1 illustrates how our method adaptively controls the size of different model components.

Method	Size	Compress	PPL
Trans. (Dai et al.)	151M	-	24.1
SRU (base)	100M	-	24.5

FAC	50M	50%	28.2
AGP	50M	50%	25.7
NP- l_0	51M	50%	26.7
FLOP -AGP	51M	50%	25.6
FLOP - l_0	50M	50%	25.3

FAC	30M	70%	31.0
AGP	30M	70%	28.4
NP- l_0	31M	70%	31.3
FLOP -AGP	31M	70%	28.1
FLOP - l_0	30M	70%	27.7

FAC	21M	80%	35.2
AGP	20M	80%	32.6
NP- l_0	18M	80%	39.1
FLOP -AGP	21M	80%	31.3
FLOP - l_0	21M	80%	31.9

Table 1: Comparison of FLOP and all baselines on the Wiki-103 dataset. We report test perplexity (PPL) at three different compression levels. All methods use adaptive embedding and softmax layers.




Figure 1: Number of parameters used in the RNN and adaptive embeddings at different compression levels. We also show the number of parameters used for the most, second most and least frequent words.

We show the overall size of recurrent encoder and adaptive embedding layers at the compression levels tested, and break down the use of parameters within three word clusters based on their frequency. FLOP learns to prune the dimension more aggressively for less-frequent words. This result showcases the benefit of adaptively reducing word dimensions.

Char-level Language Model Table 3 shows the results of pruning character-level language models. Our base model achieves a test bits-per-character score (BPC) of 1.24, which is comparable with previous reported results of RNN-based models.

As shown in Table 3, we again see the benefit

Parameters	Compression	SST2	MRPC	STS-B	QNLI	Average
125M	0%	92.43	90.9	90.22	89.77	90.83
80M	35%	92.09	88.61	88.18	89.05	89.48

Table 2: Compression on downstream fine-tuning

Method	Size	Comp.	BPC
LSTM (Wu et al.)	17M	-	1.44
QRNN (Merity et al.)	26M	-	1.33
SRU (base)	35M	-	1.24
FAC	11M	70%	1.33
AGP	11M	70%	1.27
NP- l_0	11M	70%	1.31
FLOP -AGP	11M	70%	1.27
FLOP - l_0	11M	70%	1.25
FAC	8M	80%	1.38
AGP	8M	80%	1.29
NP- l_0	8M	80%	1.34
FLOP -AGP	8M	80%	1.29
FLOP - l_0	8M	80%	1.27
FAC	4M	90%	1.47
AGP	4M	90%	1.35
NP- l_0	4M	90%	1.43
FLOP -AGP	4M	90%	1.34
FLOP - l_0	4M	90%	1.33

Table 3: Comparison of FLOP and all baselines on the Enwiki8 dataset. We report bits-per-character (BPC) on the test set. We also include previous reported results of recurrent language models on this dataset as additional data points.

Method	Size	Compress	BPC
Trans-XL (base)	41M	-	1.08
FAC	8M	80%	1.20
AGP	8M	80%	1.14
FLOP -AGP	8M	80%	1.17
FLOP - l_0	8M	80%	1.13
FLOP -AGP	4M	90%	1.25
FLOP - l_0	4M	90%	1.17

Table 4: Results of pruning Transformer-XL models on the Enwiki8 dataset. We report bits-per-character (BPC) on the test set.

of low-rank pruning, matching or improving on unstructured pruning. Furthermore, FLOP - l_0 obtains the best performance across all pruning levels. Notably, we achieve a perplexity of 1.25 at 70% compression, nearly matching the un-compressed model at 1.24.

Table 4 presents the results of pruning 12-layer

Transformer-XL models on the Enwik8 dataset. We compare FAC, unstructured AGP, FLOP-AGP and FLOP- l_0 at 80% compression level, and also report the result of FLOP variants at 90% compression. FLOP- l_0 outperforms other methods in comparison. In addition, it is able to achieve 1.17 BPC using 4M parameters, showcasing the effectiveness of our method when applied to another neural architecture.

BERT on Classification Tasks Finally, we demonstrate that our method can also be applied to language model fine-tuning on downstream tasks. We use the RoBERTa base model in this experiment. Since the model was pretrained without matrix factorization, we first compute the singular value decomposition of each matrix and then introduce the pruning mask in between the resulting factored matrices. Note that this procedure temporarily increases the total number of parameters. We compare here the final number of parameters to the initial number pre-factorization.

Our results are shown in in Table 2. We are able to conserve nearly 99% of the performance while reducing the number of parameters by 35%. Our target compression level is limited by the fact that the embedding layers consist of a significant portion of the remaining parameters. As demonstrated in the previous experiment on Wiki-103, we believe that higher levels of compression could be obtained by factorizing the embedding layer, similar to Lan et al. (2020).

7 Analysis

In this section, we perform an analysis of several aspects of our method.

Factorization One of the key hypotheses outlined in this paper is that pruning input dimensions (equivalently rows or columns of weight matrices) is a more restrictive form of pruning compared to our factorization based strategy. However, one could also argue that the factorization method works better simply because the hidden size can be initially set much larger than an unfactorized model, not because of pruning itself. For instance,

Variants	Size	0%	70%	80%	85%	90%
NP- l_0	37M	1.30	1.31 (-0.8%)	1.34 (-3.2%)	1.37 (-5.4%)	1.43 (-10.0%)
	66M	1.25	1.28 (-2.4%)	1.31 (-4.8%)	1.32 (-5.6%)	1.37 (-9.6%)
FLOP- l_0	35M	1.24	1.25 (-0.8%)	1.27 (-2.4%)	1.29 (-4.0%)	1.33 (-7.3%)

Table 5: Further comparison between factorization-based pruning FLOP and input feature pruning NP- l_0 (Louizos et al., 2018) using 6-layer SRU models and the Enwiki8 dataset. We show BPC at different compression levels and the loss of performance relative to the un-compressed model. Factorization results in less decrease in relative and absolute performance.

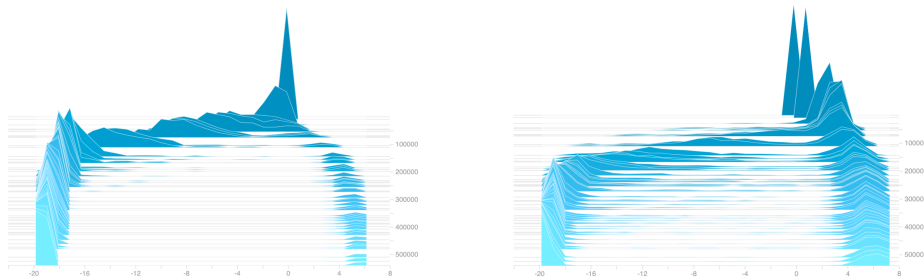


Figure 2: Histograms of HardConcrete parameters during training. We show the changes of histograms for the first SRU layer (left figure) and the last layer (right figure). We compute the histogram every 3,000 training steps.

the SRU model used by the unfactorized NP- l_0 baseline has hidden size 1536, while with factorization other baselines with a similar parameter budget use a hidden size of 3056. To avoid potential unfair comparison, we also train a large model with hidden size 2048 containing 90% more parameters, and apply the NP- l_0 baseline. This larger model obtains 1.25 BPC which is on par with the factorized base model used in previous experiments.

Table 5 compares the pruning performance of FLOP and NP- l_0 at four compression levels. We show the test BPC and the loss of performance relative to the model without pruning. These results further substantiate our hypothesis – factorization based pruning is able to retain relative model performance much more effectively than input feature pruning.

Speed analysis Thanks to its structured nature, FLOP can achieve significant computation speedup. As shown in Table 6, we achieve an inference speedup ranging from 1.5x to 2.2x for the compression levels tested, using CPUs. Similar speedups of up to 2.4x are also observed using GPUs during training. On the contrary, the computations of unstructured sparse matrices are harder to optimize. For models obtained using unstructured AGP, we experimented with the sparse matrix multiplication routine provided in Pytorch (Paszke et al., 2017) and a recent linear algebra compiler (Kjolstad et al., 2017), but were unable to achieve a speedup.

Size	Compress	Time (s)	Speedup
35M	0%	0.39	1.0x
8M	80%	0.21	1.9x
4M	90%	0.18	2.2x
41M	0%	1.33	1.0x
8M	80%	0.87	1.5x
4M	90%	0.82	1.6x

Table 6: Inference timing measurements of character-level language model using SRU (top block) and Transformer-XL (bottom block).

Learning dynamics Figure 2 demonstrates the training dynamics of the HardConcrete distribution. We plot the histogram of HardConcrete parameters α after every few thousands of training iterations. A negative value of α indicates that the associated parameter is likely to be pruned while a positive value indicates the opposite. The magnitude of the value reflects the certainty of the pruning decision. As illustrated by the figure, the distribution of α becomes bi-modal after initial exploration. Certain parameters within each layer are completely pruned while others are kept with (almost) absolute certainty. In addition, the dynamics vary across different layers. For instance, for SRU the first recurrent layer gets pruned more aggressively than the last layer.

8 Conclusion

In this work, we present a generic structured pruning method based on adaptive low-rank factorization. We systematically evaluate the performance of this method on large language models. We show that our method can provide significant speedups and compression rates on large models while losing minimal performance compared to other methods, including unstructured magnitude pruning. This work contributes to reducing the growing overhead of large language models, and shines a light on the role of model capacity in language modeling.

Acknowledgement

We would like to thank ASAPP Inc. for making this work possible. We would also like to thank Hugh Perkins, Sam Bowman, Nicholas Matthews, Josh Shapiro and the other members of the Language Technology and Research teams who helped review this work and contributed their thoughts throughout the project. We would also like to thank the EMNLP reviewers and area chair for their helpful comments.

References

- Jimmy Ba and Rich Caruana. 2014. [Do deep nets really need to be deep?](#) In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2654–2662. Curran Associates, Inc.
- Alexei Baevski and Michael Auli. 2019. [Adaptive input representations for neural language modeling](#). In *International Conference on Learning Representations*.
- Jasmijn Bastings, Wilker Aziz, and Ivan Titov. 2019. [Interpretable neural predictions with differentiable binary variables](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2963–2977, Florence, Italy. Association for Computational Linguistics.
- Shijie Cao, Chen Zhang, Zhuliang Yao, Wencong Xiao, Lanshun Nie, Dechen Zhan, Yunxin Liu, Ming Wu, and Lintao Zhang. 2019. [Efficient and effective sparse lstm on fpga with bank-balanced sparsity](#). In *Proceedings of the 2019 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays, FPGA '19*, page 63–72, New York, NY, USA. Association for Computing Machinery.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. [SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Yew Ken Chia, Sam Witteveen, and Martin Andrews. 2019. [Transformer to cnn: Label-scarce distillation for efficient text classification](#). *arXiv preprint arXiv:1909.03508*.
- Andrew M Dai and Quoc V Le. 2015. [Semi-supervised sequence learning](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 3079–3087. Curran Associates, Inc.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. [Transformer-XL: Attentive language models beyond a fixed-length context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Angela Fan, Edouard Grave, and Armand Joulin. 2020. [Reducing transformer depth on demand with structured dropout](#). In *International Conference on Learning Representations*.
- Jeffrey Flanigan, Sam Thomson, Jaime Carbonell, Chris Dyer, and Noah A. Smith. 2014. [A discriminative graph-based parser for the Abstract Meaning Representation](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1426–1436, Baltimore, Maryland. Association for Computational Linguistics.
- Jonathan Frankle and Michael Carbin. 2019. [The lottery ticket hypothesis: Finding sparse, trainable neural networks](#). In *International Conference on Learning Representations*.
- Trevor Gale, Erich Elsen, and Sara Hooker. 2019. [The state of sparsity in deep neural networks](#). *arXiv preprint arXiv:1902.09574*.
- Yunchao Gong, Liu Liu, Ming Yang, and Lubomir Bourdev. 2014. [Compressing deep convolutional networks using vector quantization](#). *arXiv preprint arXiv:1412.6115*.

- Edouard Grave, Armand Joulin, Moustapha Cissé, Hervé Jégou, et al. 2017. [Efficient softmax approximation for gpus](#). In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1302–1310. JMLR. org.
- Scott Gray, Alec Radford, and Diederik P Kingma. 2017. [Gpu kernels for block-sparse weights](#). *arXiv preprint arXiv:1711.09224*.
- Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Arda-van Pedram, Mark A Horowitz, and William J Dally. 2016. [Eie: efficient inference engine on compressed deep neural network](#). In *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*, pages 243–254. IEEE.
- Song Han, Huizi Mao, and William J Dally. 2015a. [Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding](#). *arXiv preprint arXiv:1510.00149*.
- Song Han, Jeff Pool, John Tran, and William Dally. 2015b. [Learning both weights and connections for efficient neural network](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 1135–1143. Curran Associates, Inc.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. [Distilling the knowledge in a neural network](#). *arXiv preprint arXiv:1503.02531*.
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2019. [Tinybert: Distilling bert for natural language understanding](#). *arXiv preprint arXiv:1909.10351*.
- Yoon Kim and Alexander M. Rush. 2016. [Sequence-level knowledge distillation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas. Association for Computational Linguistics.
- Fredrik Kjolstad, Shoaib Kamil, Stephen Chou, David Lugato, and Saman Amarasinghe. 2017. [The tensor algebra compiler](#). *Proc. ACM Program. Lang.*, 1(OOPSLA).
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [Albert: A lite bert for self-supervised learning of language representations](#). In *International Conference on Learning Representations*.
- Tao Lei, Yu Zhang, Sida I. Wang, Hui Dai, and Yoav Artzi. 2018. [Simple recurrent units for highly parallelizable recurrence](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4470–4481, Brussels, Belgium. Association for Computational Linguistics.
- Zhongliang Li, Raymond Kulhanek, Shaojun Wang, Yunxin Zhao, and Shuang Wu. 2018. [Slim embedding layers for recurrent neural language models](#). In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Liyuan Liu, Xiang Ren, Jingbo Shang, Xiaotao Gu, Jian Peng, and Jiawei Han. 2018. [Efficient contextualized representation: Language model pruning for sequence labeling](#). In *EMNLP*, pages 1215–1225.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Christos Louizos, Karen Ullrich, and Max Welling. 2017. [Bayesian compression for deep learning](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 3288–3298. Curran Associates, Inc.
- Christos Louizos, Max Welling, and Diederik P. Kingma. 2018. [Learning sparse neural networks through \$l_0\$ regularization](#). In *International Conference on Learning Representations*.
- André F. T. Martins, Mario A. T. Figueiredo, Pedro M. Q. Aguiar, Noah A. Smith, and Eric P. Xing. 2011. [An augmented lagrangian approach to constrained map inference](#). In *ICML*, pages 169–176.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018. [An analysis of neural language modeling at multiple scales](#). *CoRR*, abs/1803.08240.
- Dmitry Molchanov, Arsenii Ashukha, and Dmitry P. Vetrov. 2017. [Variational dropout sparsifies deep neural networks](#). In *ICML*, pages 2498–2507.
- Sharan Narang, Erich Elsen, Gregory Diamos, and Shubho Sengupta. 2017a. [Exploring sparsity in recurrent neural networks](#). *arXiv preprint arXiv:1704.05119*.
- Sharan Narang, Eric Undersander, and Gregory Diamos. 2017b. [Block-sparse recurrent neural networks](#). *arXiv preprint arXiv:1711.02782*.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. [Automatic differentiation in PyTorch](#). In *NIPS Autodiff Workshop*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving language understanding by generative pre-training](#).

- Alexander M. Rush, David Sontag, Michael Collins, and Tommi Jaakkola. 2010. [On dual decomposition and linear programming relaxations for natural language processing](#). In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1–11, Cambridge, MA. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter](#). *arXiv preprint arXiv:1910.01108*.
- Abigail See, Minh-Thang Luong, and Christopher D. Manning. 2016. [Compression of neural machine translation models via pruning](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 291–301, Berlin, Germany. Association for Computational Linguistics.
- Yuan Shangguan, Jian Li, Liang Qiao, Raziq Alvarez, and Ian McGraw. 2019. [Optimizing speech recognition for the edge](#). *arXiv preprint arXiv:1909.12408*.
- Sheng Shen, Zhen Dong, Jiayu Ye, Linjian Ma, Zhewei Yao, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. 2019. [Q-bert: Hessian based ultra low precision quantization of bert](#). *arXiv preprint arXiv:1909.05840*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. 2019. [Patient knowledge distillation for BERT model compression](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics.
- Henry Tsai, Jason Riesa, Melvin Johnson, Naveen Arivazhagan, Xin Li, and Amelia Archer. 2019. [Small and practical BERT models for sequence labeling](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics.
- Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Well-read students learn better: The impact of student initialization on knowledge distillation](#). *arXiv preprint arXiv:1908.08962*.
- Ehsan Variiani, Ananda Theertha Suresh, and Mitchel Weintraub. 2019. [West: Word encoded sequence transducers](#). In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7340–7344. IEEE.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *International Conference on Learning Representations*.
- Wei Wen, Yuxiong He, Samyam Rajbhandari, Minjia Zhang, Wenhan Wang, Fang Liu, Bin Hu, Yiran Chen, and Hai Li. 2018. [Learning intrinsic sparse structures within long short-term memory](#). In *International Conference on Learning Representations*.
- Yuhuai Wu, Saizheng Zhang, Ying Zhang, Yoshua Bengio, and Ruslan R Salakhutdinov. 2016. [On multiplicative integration with recurrent neural networks](#). In *Advances in Neural Information Processing Systems*.
- Zhuliang Yao, Shijie Cao, Wencong Xiao, Chen Zhang, and Lanshun Nie. 2019. [Balanced sparsity for efficient dnn inference on gpu](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5676–5683.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. [Ernie: Enhanced language representation with informative entities](#). *arXiv preprint arXiv:1905.07129*.
- Chenzhuo Zhu, Song Han, Huizi Mao, and William J. Dally. 2017. [Trained ternary quantization](#). In *5th International Conference on Learning Representations ICLR*.
- Michael Zhu and Suyog Gupta. 2017. [To prune, or not to prune: exploring the efficacy of pruning for model compression](#). *arXiv preprint arXiv:1710.01878*.
- Neta Zmora, Guy Jacob, Lev Zlotnik, Bar Elharar, and Gal Novik. 2019. [Neural network distiller: A python package for dnn compression research](#).

A Appendix

A.1 Optimization details

In our implementation, FLOP trains the factorized model for a number of warmup epochs and then starts pruning. Other pruning baselines use the same warmup training process, except that FAC baseline directly trains smaller factorized model from scratch. Recall our augmented Lagrangian training objective during pruning is,

$$\max_{\lambda_1, \lambda_2} \min_{\theta, \alpha} \mathbb{E}_{\mathbf{u}} \left[\frac{1}{D} \sum_{i=1}^D \mathcal{L}(\mathbf{x}_i, \mathbf{y}_i; \tilde{\theta}) \right] + g(\lambda, \alpha),$$
$$g(\lambda, \alpha) = \lambda_1 \cdot (s(\alpha) - t) + \lambda_2 \cdot (s(\alpha) - t)^2.$$

We gradually increase the target size t at a linear rate. That is, given the desired size t_{\max} , we set the sparsity at k -th pruning iteration as

$$t_k = \min\left(1, \frac{k}{m}\right) \cdot t_{\max}$$

where m is a hyperparameter specifying the number of annealing steps.

The Lagrangian multipliers are initialized to zero at the start of training. We perform joint gradient updates for the parameters and Lagrangian multipliers at every iteration, but use and tune a different learning rate for Lagrangian multipliers. For each training batch, we sample the pruning mask $\mathbf{z} = \{z_1, \dots, z_n\}$ and share it across the training examples within the batch. Since the pruning mask is shared, we can select parameters that are only active for the current batch and compute smaller matrix multiplications in forward and backward passes. This can result in training speedup when \mathbf{z} becomes sparse.

A.2 Experimental Details

Our experiments are performed using the standard train/dev/test splits of Wiki-103, Enwik8 and GLUE benchmarks. We describe training configurations in the following paragraphs. Detailed experimental setup can be found at <https://github.com/asappresearch/flop>.

SRU Following the practice of [Lei et al. \(2018\)](#), for the Enwik8 dataset we train a 6-layer SRU model using a batch size of 64 and an unroll length of 256. We use a hidden size of 3056 and set the initial factorization dimension r of the parameter matrices to 512. That is, we replace each weight matrix \mathbf{W} in SRU using an explicit factorization

PQ with an inner dimension of 512. We train the model without pruning for 30 warmup epochs, and start pruning for a maximum of 100 epochs.

For the Wiki-103 dataset, our 12-layer SRU base model uses a hidden dimension of 2048 and a factorization dimension of 512 for weight matrices in SRU. Following [Baevski and Auli \(2019\)](#), the adaptive embedding layer uses 1024, 256 and 64 dimensions respectively for the 20K most frequent words, 40K less frequent words and the rest least frequent words. We train 50 warm-up epochs and start the pruning process for an addition of 100 epochs. We use a batch size of 64 or 96 and an unroll length of 256.

For all SRU runs, we use inverse-square-root learning rate scheduling ([Vaswani et al., 2017](#)) and a learning rate of $\frac{l_0}{\sqrt{d}}$ where d is the hidden size and l_0 is the initial factor. We set $l_0 \in \{2, 3\}$ for model parameters. For AGP methods, we tune the start and end epoch of the compression scheduler. For l_0 regularization, we tune the learning rate $l_0 \in \{3, \dots, 6\}$ for Lagrangian multipliers.

Transformer-XL Following [Dai et al. \(2019\)](#), we train Transformer-XL base model using cosine learning rate scheduling. For the 12-layer base model, we train a maximum of 200k iterations, a batch size of 48 and an initial learning rate of 0.0003 and use 8 GPUs in parallel. For pruning runs, we train up to 300k iterations using a learning rate of 0.00025, a batch size of 32 and 4 GPUs in parallel for each run. We use the same inverse-square-root learning rate scheduling for Lagrangian multipliers and set $l_0 \in \{0.5, 1.0, 1.5\}$.