

Pruning Pre-trained Language Models with Principled Importance and Self-regularization

Siyu Ren Kenny Q. Zhu*

Shanghai Jiao Tong University

Shanghai, China

roy0702@sjtu.edu.cn, kzhu@cs.sjtu.edu.cn

Abstract

Iterative pruning is one of the most effective compression methods for pre-trained language models. We discovered that finding the optimal pruning decision is an equality-constrained 0-1 Integer Linear Programming problem. The solution to this optimization problem leads to a principled importance criterion which we use to rank parameters during iterative model pruning. To mitigate the poor generalization at high sparsity levels, we propose a self-regularization scheme where model prediction is regularized by the latest checkpoint with increasing sparsity throughout pruning. Our experiments on natural language understanding, question answering, named entity recognition, and data-to-text generation with various Transformer-based PLMs show the effectiveness of the approach at various sparsity levels.

1 Introduction

Pre-trained language models (PLMs) (Devlin et al., 2019; Radford et al., 2018) have significantly advanced the state-of-the-art in various natural language processing tasks (Wang et al., 2018; Zhou and Lampouras, 2020; Dušek et al., 2020; Radev et al., 2020). However, these models often contain a vast amount of parameters, posing non-trivial requirements for storage and computation. Due to this inefficiency, the applications of PLMs in resource-constrained scenarios are still limited.

To resolve the above challenge, model compression (Sun et al., 2019; Ben Noach and Goldberg, 2020; Lan et al., 2020) has been actively studied to make PLMs meet the practical requirement. Among them, iterative pruning methods are widely adopted at only a tiny expense of model performance when adapting PLMs to downstream tasks. During the course of iterative pruning, model parameters can not only be updated but also

be pruned based on the rank of their importance scores in order to satisfy the cardinality constraint. Prevalent importance criteria are based on the parameter’s magnitude (Zhu and Gupta, 2017; Renda et al., 2020) or sensitivity (Louizos et al., 2018; Sanh et al., 2020; Liang et al., 2021; Zhang et al., 2022). Parameters with low importance scores are pruned and are expected to have little impact on model performance.

Despite the empirical success, existing importance criteria for model pruning still face two major limitations: (1) they are heuristically defined and may not accurately quantify a parameter’s contribution to the learning process, e.g., absolute weight value in magnitude-based pruning and gradient-weight product in sensitivity-based pruning; (2) they determine the importance of each parameter individually without considering the effect of coinstantaneous parameter updates on model performance, e.g., sensitivity is estimated by the absolute change in training error if only a single parameter is pruned and others remain unchanged.

In this paper, we rethink the design of the importance criterion for model pruning from an optimization perspective. We begin by analyzing the temporal variation of any given learning objective based on a single-step gradient descent update under the iterative pruning setting. We show that finding the optimal pruning decision can be framed as solving an equality-constrained 0-1 Integer Linear Programming (ILP) problem, where the constraint is defined by the specified sparsity. The resulting problem is a particular case of a general 0-1 Knapsack problem in which the weight for each item is the same. The solution to this problem naturally leads to a principled importance criterion which we use to rank all model parameters and derive the optimal stepwise pruning decision.

When a high sparsity (e.g., 80%~90%) is pursued, the limited capacity often renders the pruned model fails to retain satisfactory performance with

* The corresponding author.

conventional fine-tuning. To further improve the model’s generalization ability, we propose a self-regularization scheme, where the model prediction is regularized by the latest best-performing model checkpoint during pruning. We show that such a scheme eases model learning with decreasing capacity and effectively yields a tighter upper bound of expected generalization error than learning from training data alone.

To validate the effectiveness of our approach, dubbed PINS (Pruning with principled Importance and Self-regularization), we conducted extensive experiments with various pre-trained language models on a wide variety of tasks, including natural language understanding on GLUE (Wang et al., 2018)), question answering on SQuAD (Rajpurkar et al., 2016), named entity recognition on CoNLL 2003 (Tjong Kim Sang and De Meulder, 2003), and data-to-text generation on WebNLG (Zhou and Lampouras, 2020), DART (Radev et al., 2020), and E2E (Dušek et al., 2020). Experimental results show that PINS provides more accurate models at different sparsity levels. Detailed analysis shed further light on some intriguing properties of models pruned by PINS. By exploiting the resulting high sparsity, we show that the storage/inference can be reduced/accelerated by 8.9x and 2.7x using CSR format and a sparsity-aware inference runtime (Kurtz et al., 2020) on consumer-level CPUs¹.

In summary, our contributions are:

- We establish the equivalence between the optimal pruning decision and the solution to an equality-constrained 0-1 Integer Linear Programming problem. The solution to this problem leads to a principled importance criterion that can be used to rank parameters during iterative pruning.
- We propose a simple yet effective self-regularization scheme to enhance the model’s generalization capability, especially under a high-sparsity regime.
- Comprehensive experiments and analyses confirm the effectiveness of our approach at various sparsity levels.

¹Code available at <https://github.com/DRSY/PINS>

2 Background and Related Work

In this section, we review the necessary background on Transformer-based pre-trained language models and popular importance criteria for iterative pruning.

2.1 Transformer-based Pre-trained Language Models

Most existing pre-trained neural language models (Radford et al., 2018; Devlin et al., 2019; Wang et al., 2020; Clark et al., 2020) are based on the Transformer (Vaswani et al., 2017) architecture, which consists of several identical blocks of self-attention and feedforward network. After pre-training on a massive amount of unlabeled general-domain corpus in a self-supervised learning manner, these models exhibit superior performance on various downstream tasks via fine-tuning. However, good generalization performance comes at the cost of a vast amount of parameters. For example, the base version of BERT has 110M parameters and leads to more than 400MB of disk storage. Therefore, how to effectively reduce model size while preserving as much task accuracy as possible remains a challenging research problem.

2.2 Iterative Pruning

Pruning methods can be divided into two categories: one-shot pruning (Lee et al., 2018; Frankle and Carbin, 2018) and iterative pruning (Louizos et al., 2018; Sanh et al., 2020; Zhang et al., 2022). One-shot pruning removes parameters of low importance after training. It is efficient but ignores the complicated training dynamics when applied to modern large neural language models. On the contrary, iterative pruning performs training and pruning simultaneously. Therefore, the resulting sparsity pattern is aware of the complex dynamics of parameters through the course of training and delivers considerable improvement compared to one-shot pruning.

Let $\theta^{(t)} = \{\theta_1^{(t)}, \theta_2^{(t)}, \dots, \theta_d^{(t)}\}$ denote the d -dimensional model parameters at t -th training iteration, the typical updating rule of iterative pruning can be formulated as:

$$\hat{\theta}^{(t+1)} = \theta^{(t)} - \eta^{(t)} \nabla_{\theta} \mathcal{L}(\theta^{(t)}) \quad (1)$$

$$\theta^{(t+1)} = \hat{\theta}^{(t+1)} \odot M^{(t)} \quad (2)$$

where $\eta^{(t)}$ is the learning rate at time step t and \mathcal{L} is the learning objective. The temporarily updated

$\hat{\theta}^{(t+1)}$ is further pruned by the binary mask $M^{(t)} \in \{0, 1\}^d$, which is computed based on a given importance criterion $S^{(t)}$:

$$M_i^{(t)} = \begin{cases} 1, & \text{if } S_i^{(t)} \text{ is in the top-}r^{(t)} \text{ of } S^{(t)} \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where $r^{(t)} \leq d$ indicates the number of remaining parameters at time step t according to a given sparsity scheduler.

2.3 Importance Criteria for Model Pruning

Popular importance criteria for model pruning include parameters' magnitude and sensitivity.

Magnitude is a simple yet effective importance criterion that is widely used for model pruning. It estimates the importance of each parameter as its absolute value, i.e., $S_i^{(t)} = |\theta_i^{(t)}|$. Despite its simplicity, the magnitude cannot accurately gauge the importance of parameters because even parameters with small magnitude can have a large impact on the model prediction due to the complex compositional structure of PLMs.

Sensitivity is another useful importance criterion. It estimates the importance of each parameter as the absolute change of the learning objective if the parameter is pruned, i.e., set to zero. The mathematical formulation of the sensitivity of i -th parameter is given by:

$$S_i^{(t)} = |\mathcal{L}(\theta_{-i}^{(t)}) - \mathcal{L}(\theta^{(t)})| \quad (4)$$

$$\approx |g_i^{(t)} \theta_i^{(t)}| \quad (5)$$

where $\theta_{-i}^{(t)}$ is identical to $\theta^{(t)}$ except that the i -th entry is set to zero and $g_i^{(t)}$ is the gradient of i -th entry. Though taking the training dynamics into account, sensitivity still estimates the importance of each parameter individually without considering the effect of holistic parameter update.

3 Methodology

Instead of heuristically defining the importance criterion as in prior pruning methods, we take a step back and rethink the design of the importance criterion for model pruning from an optimization perspective. From our analysis, we draw an equivalence between finding the optimal stepwise pruning decision and solving an equality-constrained

0-1 Integer Linear Programming problem. We further show that the optimal solution to this problem leads to a new importance criterion for model pruning. Moreover, we propose a simple yet effective self-regularization scheme to facilitate the generalization ability of the sparse model. We elucidate our analysis in Section 3.1 and describe our self-regularization scheme in Section 3.2.

3.1 Rethinking Importance Criterion from the Optimization Perspective

Without loss of generality, we denote \mathcal{L} as the learning objective when adapting a pre-trained language model f with parameter θ to a downstream task. At t -th training iteration, we denote the current model parameters as $\theta^{(t)}$ and the evaluated learning objective as $\mathcal{L}(\theta^{(t)})$.

The temporal variation of the learning objective $\mathcal{L}(\theta^{(t)})$ at time step t is given by the second-order Taylor series expansion:

$$\Delta\mathcal{L}^{(t)} = \mathcal{L}(\theta^{(t)} + \Delta\theta^{(t)}) - \mathcal{L}(\theta^{(t)}) \quad (6)$$

$$= \nabla_{\theta} \mathcal{L}(\theta^{(t)})^{\top} \Delta\theta^{(t)} +$$

$$\frac{1}{2} \Delta\theta^{(t)\top} \mathbf{H}^{(t)} \Delta\theta^{(t)} + o(|\Delta\theta^{(t)}|^2) \quad (7)$$

where $\mathbf{H}^{(t)}$ is the Hessian matrix at step t . It is known that the largest eigenvalue λ_{max} of Hessian matrices in a PLM is typically small (Shen et al., 2019), i.e., $\Delta\theta^{(t)\top} \mathbf{H}^{(t)} \Delta\theta^{(t)} \leq \lambda_{max} |\Delta\theta^{(t)}|_2^2 \approx 0$. Thus, we ignore the second-order term as well as the infinitesimal of higher order in Eq. (7):

$$\Delta\mathcal{L}^{(t)} = \nabla_{\theta} \mathcal{L}(\theta^{(t)})^{\top} \Delta\theta^{(t)}$$

$$= \sum_{i=1}^d g_i^{(t)} \cdot \Delta\theta_i^{(t)} \quad (8)$$

Under the iterative pruning setting, the actual temporal variation $\Delta\theta_i^{(t)}$ of i -th parameter depends on whether it is allowed to be updated or forced to zeroed out. Formally, we use a binary variable $x_i^{(t)}$ to indicate the pruning decision of i -th parameter at time step t , i.e., $x_i^{(t)} = 1$ means $\theta_i^{(t)}$ is updated and $x_i^{(t)} = 0$ means $\theta_i^{(t)}$ is pruned. The temporal variation in Eq. (8) can now be rewritten as:

$$\Delta\mathcal{L}^{(t)} = \sum_{i=1}^d g_i^{(t)} (x_i^{(t)} \Delta\hat{\theta}_i^{(t)} + (1 - x_i^{(t)}) (-\theta_i^{(t)})) \quad (9)$$

where $\Delta\hat{\theta}_i^{(t)} = -\eta^{(t)} g_i^{(t)}$ is the gradient descent update. Finding the optimal pruning decision that

leads to the smallest $\Delta\mathcal{L}^{(t)}$ is now converted to an equality-constrained 0-1 integer linear programming (ILP) problem of variables $\mathbf{x}^{(t)}$:

$$\begin{aligned} \tilde{\mathbf{x}}^{(t)} &= \arg \min_{\mathbf{x}^{(t)}} \Delta\mathcal{L}^{(t)} \\ \text{s.t. } \sum_{i=1}^d \mathbf{x}_i^{(t)} &= r^{(t)}, \mathbf{x}_i^{(t)} \in \{0, 1\} \end{aligned} \quad (10)$$

where $r^{(t)}$ is the number of remaining parameters at step t according to the pre-defined sparsity scheduler. If we consider each parameter $\theta_i^{(t)}$ as an item and $r^{(t)}$ as the total capacity, the problem that Eq. (10) defines can be treated as a special case of 0-1 Knapsack problem where the weight for each item is one and the value for each item is given by:

$$\mathbf{S}_i^{(t)} = -\mathbf{g}_i^{(t)} \Delta\hat{\theta}_i^{(t)} - \mathbf{g}_i^{(t)} \theta_i^{(t)} \quad (11)$$

Contrary to the general 0-1 Knapsack problem which is known to be NP-complete, fortunately, the equal-weight 0-1 Knapsack is a P problem. Its optimal solution can be obtained by sorting items in descending order according to their values and selecting the top- $r^{(t)}$ ones:

$$\tilde{\mathbf{x}}_i^{(t)} = \begin{cases} 1, & \text{if } \mathbf{S}_i^{(t)} \text{ is in the top-}r^{(t)} \text{ of } \mathbf{S}^{(t)} \\ 0, & \text{otherwise} \end{cases} \quad (12)$$

Putting it in the context of iterative pruning, our analysis theoretically reveals the validity of: (1) selecting parameters based on the ranking of certain importance criterion; (2) using Eq. (11) as a principled new importance criterion.

3.2 Self-regularization

In vanilla fine-tuning, the learning objective \mathcal{L} is defined as the training error \mathcal{L}_{er} (a.k.a empirical risk in statistical learning) over the empirical data distribution. However, minimizing such training error does not translate to good generalization. Moreover, as iterative pruning proceeds, the number of non-zero parameters in the model monotonically decreases. The reduced model capacity increases the learning difficulty (Lopez-Paz et al., 2015; Mirzadeh et al., 2019) and usually leads to degenerated generalization performance of the sparsified model (Sanh et al., 2020).

Confronting the above challenges, we propose an effective self-regularization scheme tailored to

improving the model’s generalization ability during iterative pruning. Concretely, besides learning from the hard label of training data, the output of the current model with parameter $\theta^{(t)}$ is also regularized by the output of the latest best-performing model checkpoint with parameter $\theta^{(t_l)}$, where $t_l \leq t$ denotes the time step at which the latest checkpoint was saved. The learning objective of self-regularization is defined as:

$$\mathcal{L}_{sr} = \mathcal{D}(y_{\theta^{(t)}}, y_{\theta^{(t_l)}}) \quad (13)$$

where \mathcal{D} can be any divergence metric, e.g., KL-divergence for classification tasks. \mathcal{L}_{sr} is then integrated with the original learning objective, i.e., $\mathcal{L} = \mathcal{L}_{er} + \mathcal{L}_{sr}$.

Why does self-regularization work? Our self-regularization is similar to teacher-student knowledge distillation in the sense that the model output is regularized by the output of another model. However, the most critical difference is that the “teacher” in self-regularization is instantiated by checkpoint with increasing sparsity, such that the capacity gap between “teacher” and “student” is dynamically adjusted. We theoretically justify the effectiveness of self-regularization as follows:

Theorem 1. *Let t_i and t_j where $t_i \geq t_j$ denote the time steps at which two different checkpoints are saved; Let $R(f_{\theta^{(t \leftarrow t_i)}})$ and $R(f_{\theta^{(t \leftarrow t_j)}})$ denote the expected generalization error of models learned from $f_{\theta^{(t_i)}}$ and $f_{\theta^{(t_j)}}$; Let n denotes the size of training data; $|\cdot|_C$ denotes a capacity measure of function class \mathcal{F}_θ . Based on previous expositions on VC theory (Vapnik, 1998), we have the following asymptotic generalization bounds hold:*

$$\begin{aligned} R(f_{\theta^{(t \leftarrow t_i)}}) &\leq \underbrace{O\left(\frac{|\mathcal{F}_{\theta^{(t)}}|_C}{n^{\alpha_i}}\right) + \inf_{\mathcal{F}_{\theta^{(t \leftarrow t_i)}}} R(f_{\theta^{(t)}})}_{\text{bound}(f_{\theta^{(t \leftarrow t_i)}})} \\ R(f_{\theta^{(t \leftarrow t_j)}}) &\leq \underbrace{O\left(\frac{|\mathcal{F}_{\theta^{(t)}}|_C}{n^{\alpha_j}}\right) + \inf_{\mathcal{F}_{\theta^{(t \leftarrow t_j)}}} R(f_{\theta^{(t)}})}_{\text{bound}(f_{\theta^{(t \leftarrow t_j)}})} \end{aligned}$$

Because $\theta^{(t_i)}$ is a later checkpoint with higher sparsity than $\theta^{(t_j)}$, we have the learning speed $1 \geq \alpha_i \geq \alpha_j \geq \frac{1}{2}$, then the following inequality holds with high probability:

$$\text{bound}(f_{\theta^{(t \leftarrow t_i)}}) \leq \text{bound}(f_{\theta^{(t \leftarrow t_j)}})$$

In summary, self-regularization works by enabling a tighter generalization bound compared to learning from training data alone or a static dense teacher as in knowledge distillation. Please refer to Appendix B for detailed derivation.

3.3 The Algorithm

Here we formally summarize our algorithm PINS (Pruning with principled Importance and Self-regularization) in Algorithm 1:

Algorithm 1 PINS

Input: Training set $\mathcal{D}_{tr} = \{(x_i, y_i)\}_{i=1}^N$; Validation set \mathcal{D}_{val} ; pre-trained parameters θ ; maximum training steps T ; evaluation interval t_{eval} .

Initialize: $\theta^{(0)} \leftarrow \theta$, $t_l \leftarrow 0$, best validation accuracy $acc_{t_l} \leftarrow -\text{INF}$.

- 1: **for** $t = 0$ to $T - 1$ **do**
- 2: Sample a mini-batch (x, y) from \mathcal{D}_{tr}
- 3: Compute current model’s output $y_{\theta^{(t)}}$
- 4: Compute latest best-performing checkpoint’s output $y_{\theta^{(t_l)}}$
- 5: Compute \mathcal{L} based on $y_{\theta^{(t)}}$, $y_{\theta^{(t_l)}}$ and y
- 6: Compute $\mathcal{S}^{(t)}$ via Eq. (11)
- 7: Compute $\theta^{(t+1)}$ via Eq. (2) and Eq. (3)
- 8: **if** $t \% t_{eval} = 0$ and $acc_t > acc_{t_l}$ **then**
- 9: $acc_{t_l} \leftarrow acc_t$, $\theta^{(t_l)} \leftarrow \theta^{(t)}$

Output: the pruned parameters $\theta^{(T)}$.

4 Experiments

In this section, We compare PINS with state-of-the-art pruning algorithms and perform detailed analysis to understand the effectiveness of PINS.

4.1 Setup

4.1.1 Tasks

We conduct experiments on a comprehensive spectrum of tasks following standard data splits.

Natural Language Understanding. We opt for tasks from the GLUE (Wang et al., 2018) benchmark, including linguistic acceptability (CoLA), natural language inference (RTE, QNLI, MNLI), paraphrase (MRPC, QQP), sentiment analysis (SST-2) and textual similarity (STS-B). Because the official test set of GLUE is hidden, we randomly split a small portion of training set as validation set and treat the original validation set as test set.

Question Answering. We use SQuAD v1.1 (Rajpurkar et al., 2016) as a representative dataset for

extractive question answering following previous work (Zhang et al., 2022).

Named Entity Recognition. We also examine our approach on CoNLL 2003 (Tjong Kim Sang and De Meulder, 2003) for token-level named entity recognition task.

Data-to-Text Generation. Besides language understanding tasks, we also extend our evaluation to data-to-text generation on three datasets: E2E (Dušek et al., 2020), DART (Radev et al., 2020), and WebNLG (Zhou and Lampouras, 2020), which involves generating a piece of fluent text from a set of structured relational triples.

4.1.2 Baselines

Magnitude-based. Iterative magnitude pruning (IMP) (Zhu and Gupta, 2017) is the state-of-the-art magnitude-based approach.

Sensitivity-based. l_0 -regularization (Louizos et al., 2018) trains masking variables via re-parametrization trick with l_0 penalty; SMvP (Sanh et al., 2020) uses accumulated sensitivity as importance metric; PST (Li et al., 2022) proposed a hybrid importance criterion combining both magnitude and sensitivity; PLATON (Zhang et al., 2022) uses a modified variant of sensitivity by exponential moving average and uncertainty re-weighting.

4.1.3 Implementation Details

We mainly conduct experiments on the pre-trained BERT_{base} (Devlin et al., 2019) as a pruning target for all tasks except data-to-text generation. We defer the pruning results of MiniLM_{12L-384H} (Wang et al., 2020) and Electra_{base} (Clark et al., 2020) to Appendix A. For data-to-text generation, we adopt the pre-trained GPT-2 (Radford et al., 2018) following a prior study (Li et al., 2022).

During pruning, we employ the cubic sparsity scheduler (Sanh et al., 2020; Zhang et al., 2022) to gradually increase the sparsity level from 0 to the specified target sparsity. To avoid tremendous computation cost brought by hyper-parameter tuning, we only search the batch size from $\{16, 32\}$ and fix the learning rate as $3e-5$ for all experiments on GLUE and CoNLL. For SQuAD v1.1, we fix the batch size as 16 and the learning rate as $3e-5$ following Zhang et al. (2022). We adopt AdamW (Loshchilov and Hutter, 2017) as the default optimizer. To reduce the variance induced by mini-batch sampling, we adopt a smoothing technique similar to PLATON. We run each experi-

Sparsity	Method	RTE Acc	MRPC F1	STS-B Pearson	CoLA Mcc	SST-2 Acc	QNLI Acc	MNLI Acc	QQP Acc	Avg.
0%	Fine-tune†	69.3	90.3	90.2	58.3	92.4	91.3	84.0	91.5	83.4
80%	IMP†	65.7	86.2	86.8	42.5	84.3	89.2	82.2	86.0	77.9
	l_0 -regularization†	63.2	80.2	82.8	0.0	85.0	85.0	80.8	88.5	70.7
	SMvP†	62.8	86.7	87.8	48.5	89.0	88.3	81.9	90.6	79.5
	PST	63.0	87.4	88.0	44.6	89.3	88.3	79.3	88.9	78.6
	PLATON†	68.6	89.8	89.0	54.5	91.2	90.1	83.3	90.7	82.2
	PINS (ours)	72.7	90.9	89.2	57.1	91.9	91.2	83.9	90.9	83.5
90%	IMP†	57.4	80.3	83.4	18.3	80.7	86.6	78.9	78.8	70.5
	l_0 -regularization†	59.9	79.5	82.7	0.0	82.5	82.8	78.4	87.6	69.1
	SMvP†	58.8	85.9	86.5	0.0	87.4	86.6	80.9	90.2	72.1
	PST‡	62.8	85.6	81.7	42.5	88.7	86.0	76.7	83.9	76.0
	PLATON†	65.3	88.8	87.4	44.3	90.5	88.9	81.8	90.2	79.6
	PINS (ours)	68.5	90.1	87.9	49.8	91.0	89.5	82.7	90.6	81.3

Table 1: Results with BERT_{base} on the GLUE development set. For MNLI, the results are averaged on MNLI-m and MNLI-mm. † indicates the results are directly quoted from Zhang et al. (2022) while ‡ indicates the results are reported by Li et al. (2022).

Sparsity	80%	70%	60%	50%
Fine-tune†	88.1			
IMP†	82.9	86.5	86.7	87.0
l_0 -regularization†	81.9	82.8	83.9	84.6
SMvP†	–	84.6	–	85.8
PLATON†	86.1	86.7	86.9	87.2
PINS (ours)	86.4	86.9	87.4	88.0

Table 2: Results with BERT_{base} on SQuAD v1.1. † indicates numbers reported from Zhang et al. (2022). F1 score is reported as evaluation metric.

ment five times with different random seeds and report the average results (significance tests with p -value < 0.05 are conducted for all performance gains).

4.2 Main Results

4.2.1 Comparison with Baselines

Natural language understanding We present the experimental results on GLUE at high sparsity, i.e., 80% and 90% in Table 1. Among all baselines, sensitivity-based methods generally achieve better results than magnitude-based IMP, which implies the importance of training dynamics when designing pruning criteria. We can see that PINS delivers more accurate sparsified models on all datasets at both sparsity levels. The advantage of PINS is more evident on small datasets. For example, PINS outperforms the previous best-performing baseline (PLATON) by 4.1 and 2.6 points on RTE and CoLA at 80% sparsity, where there are only a few thousand training data. Under extremely high sparsity, i.e., 90%, PINS is still

Sparsity	Method	P	R	F1
0%	Fine-tune	93.5	94.6	94.0
70%	IMP	90.7	91.8	91.2
	SMvP	92.9	94.1	93.5
	PINS(ours)	93.5	94.3	93.9
80%	IMP	84.4	87.3	85.8
	SMvP	92.1	93.1	92.6
	PINS(ours)	92.8	93.8	93.3

Table 3: Results with BERT_{base} on CoNLL 2003. P and R stands for Precision and Recall respectively.

able to retain 97.5% overall performance of fine-tuning, outperforming 95.4% of the previous best method PLATON. Notably, PINS even surpasses fine-tuning on RTE and MRPC at 80% sparsity. This can be attributed to the fact that PLMs are heavily over-parameterized and PINS can effectively identify parameters crucial to the task to realize low bias and low variance simultaneously.

Question answering Table 2 summarizes the pruning results on SQuAD v1.1. Interestingly, IMP outperforms all sensitivity-based methods except for PLATON at all considered sparsity levels, in contrast to the observations on GLUE. Our method, however, consistently yields the best performance at all sparsity settings.

Named entity recognition Table 3 demonstrates the pruning results on CoNLL 2003 dataset for named entity recognition. At 70% sparsity, our method almost matches the performance of fine-tuning, outperforming baselines on all evaluation metrics. The gain of PINS is more prominent

Sparsity	Method	E2E			DART		WebNLG	
		BLEU	ROUGE-L	METEOR	BLEU	BLEURT	BLEU	BLEURT
0%	Fine-tune	69.4	71.1	46.2	46.6	0.30	46.9	0.23
80%	IMP	69.3	71.0	45.8	44.9	0.22	39.9	0.00
	PST	69.4	70.8	45.9	44.1	0.22	44.3	0.16
	PINS (ours)	69.6	71.8	46.6	46.2	0.29	45.5	0.18

Table 4: Results with GPT-2 on data-to-text generation datasets. The higher the BLEU, ROUGE-L, METEOR, and BLEURT scores are, the better the performance.

Sparsity	Method	RTE	MRPC	STS-B	CoLA	SST-2	QNLI	MNLI	QQP	Avg.
		Acc	F1	Pearson	Mcc	Acc	Acc	Acc	Acc	
0%	Fine-tune	69.3	90.3	90.2	58.3	92.4	91.3	84.0	91.5	83.4
50%	PINS	70.8	91.4	89.7	60.6	92.9	91.8	85.1	91.3	84.2
30%	PINS	71.7	91.2	89.8	60.4	93.3	92.0	85.1	91.5	84.4

Table 5: Results with BERT_{base} on the GLUE development set under medium-to-low sparsity regime. Numbers are the mean of five trials with different random seeds. PINS outperforms fine-tuning at medium-to-low sparsity.

when further increasing sparsity.

Data-to-text generation Table 4 shows the pruning results on E2E, DART and WebNLG at 80% sparsity. PINS achieves the best performance on all three datasets in all evaluation metrics. In particular, PINS delivers performance even better than fine-tuning on the E2E dataset by 0.7 ROUGE-L and 0.4 METEOR scores, respectively. We posit that this is due to the relative easiness of E2E compared to the other two datasets.

4.2.2 Results at Medium-to-Low Sparsity

The typical utility of pruning is to produce a sparse yet competitive model that can benefit downstream applications in terms of efficiency without sacrificing much task accuracy. We hypothesize that PINS might also bring a regularization effect compared to vanilla fine-tuning under the medium-to-low sparsity regime.

As shown in Table 5, when specifying a medium-to-low sparsity, e.g., 50%~30%, our method can effectively play a role of regularization and improve model performance compared to vanilla fine-tuning. With half of the parameters being pruned, the sparse model produced by PINS outperforms fine-tuning by 1 percentage point on the GLUE score. This observation suggests that appropriate pruning can effectively reduce variance without hurting model expressiveness.

4.3 Ablation Study

The self-regularization scheme is proposed and integrated into PINS to improve model generaliza-

\mathcal{L}	RTE	CoLA	MRPC
empirical risk	70.9	55.4	90.6
w/ knowledge distillation	70.3	56.0	90.6
w/ self-regularization	72.7	57.1	90.9

Table 6: Ablation Study with BERT_{base} on the learning objective during iterative pruning at 80% sparsity.

tion. Here we investigate the effectiveness of self-regularization by comparing it to the conventional knowledge distillation scheme and the classical empirical risk minimization scheme.

The pruning results of using the three different learning objectives on RTE, CoLA, and MRPC are listed in Table 6. Pruning with PINS using classical empirical risk minimization still achieves performance better than existing baselines (Table 1). Learning from a densely fine-tuned BERT_{base} as the teacher does not always improve and sometime may even hurt performance. In contrast, our proposed self-regularization consistently boosts model performance, which echoes our theoretical justification in Section 3.2.

4.4 Analysis

We provide an in-depth analysis of various importance criteria to uncover more valuable insights.

Sparsity pattern of weight matrices We are interested in the sparsity pattern produced by different pruning criteria. To this end, we plot the remaining parameters' distribution of the same weight matrix in BERT_{base} pruned via magnitude, sensitivity, and PINS in Figure 1. We observe

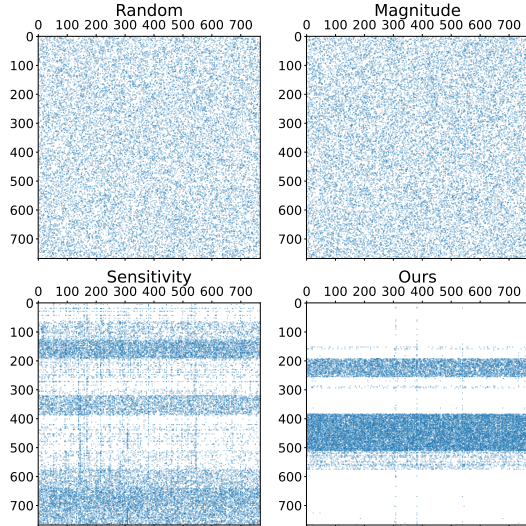


Figure 1: Sparsity pattern (80%) of the same weight matrix in $BERT_{base}$ trained on SST-2. See Appendix C for more details on the matrix rank distribution.

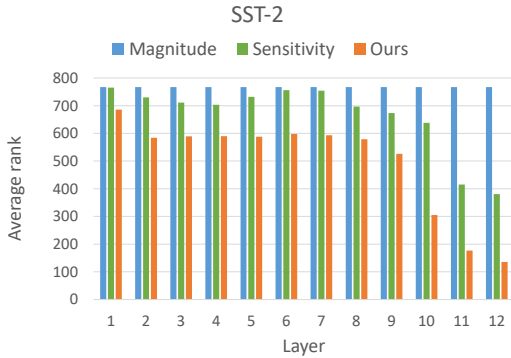


Figure 2: Layerwise distribution of average matrix rank in $BERT_{base}$ pruned at 80% sparsity on SST-2.

that magnitude-based pruning generates a sparsity pattern close to randomness. Sensitivity-based pruning produces a more structured pattern where the remaining parameters tend to occupy complete rows. Interestingly, the sparsity pattern produced by PINS exhibits the highest concentration on specific rows. This implies that the parameters contributing most to the end-task are preferably distributed in a structured way and PINS is more effective at extracting such patterns.

Layerwise rank distribution The highly structured sparsity pattern generated by PINS intrigues our interest to further analyze the intrinsic property of parameter matrices after pruning. Specifically, we inspect the matrix rank as it is usually associated with the complexity of matrix. To this end, we visualize the layerwise rank distribution of $BERT_{base}$ pruned using different importance

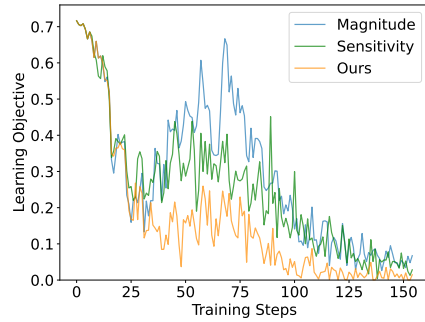


Figure 3: Change of learning objective (cross-entropy) during iterative pruning on SST-2.

Sparsity	Time(s)	Storage(MB)	Acc.
0%	0.110 (1.0x)	340 (1.0x)	69.3
80%	0.041 (2.7x)	38 (8.9x)	69.0

Table 7: Practical time and storage efficiency gain on RTE with DeepSparse and CSR format. Inference is performed on Intel Xeon E5-2640 CPU with batch size 1.

criteria on SST-2 dataset. As shown in Figure 4, magnitude pruning produces sparse matrices that are still near full-rank despite containing 80% zeros. Sensitivity pruning tends to generate sparsity pattern with lower rank compared to magnitude pruning. Notably, model pruned by PINS shows consistently lower matrix rank than the other two criteria. This implies that PINS is more effective at identifying the low-dimensional task representation during adaptation, which is usually correlated with tighter generalization bounds (Arora et al., 2018; Aghajanyan et al., 2021).

Empirical validation of importance criterion

In Section 3.1 we prove that the pruning decision derived by our importance criterion is theoretically optimal. Here we empirically validate this point by visualizing the change of learning objective as pruning proceeds. Figure 3 illustrates that our importance criterion indeed leads to the most significant decrease in the learning objective compared to heuristical ones like magnitude and sensitivity.

4.5 Efficiency Gain

We can exploit the resulting high sparsity to attain practical efficiency gain on storage and inference speed. We first apply quantization upon the pruned model and transform it into INT8 data type before saving it using Compressed Sparse Row (CSR) format. We then leverage a sparsity-aware runtime (Kurtz et al., 2020) for accelerating inference. As shown in Table 7, on the RTE dataset, the disk

space and inference time of BERT_{base} pruned at 80% sparsity can be reduced by roughly 8.9x and 2.7x respectively with negligible accuracy loss.

5 Conclusion

We present PINS, a new iterative pruning method that hinges on a principled weight importance criterion to deliver the optimal stepwise pruning decision. Integrated with a self-regularization scheme tailored to pruning-during-adaptation, PINS allows for provably better generalization ability. Empirical experiments and analyses confirm the effectiveness of our method and shed further light on the different sparsity patterns produced by PINS and other existing methods.

Limitations

Compared to the empirical risk minimization scheme, the introduced self-regularization scheme incurs certain overhead because each mini-batch of data will go through two models. For BERT_{base} scale pre-trained language models, the additional memory overhead is about 27% and the additional training time overhead is about 30%. Nevertheless, once pruned, the sparsified model can enjoy considerable efficiency gains in terms of storage and inference time. Therefore, this is a trade-off that future practitioners might need to consider.

Acknowledgments

This work was generously supported by the CMB Credit Card Center & SJTU joint research grant, and Meituan-SJTU joint research grant.

References

Armen Aghajanyan, Sonal Gupta, and Luke Zettlemoyer. 2021. [Intrinsic dimensionality explains the effectiveness of language model fine-tuning](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7319–7328, Online. Association for Computational Linguistics.

Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. 2018. Stronger generalization bounds for deep nets via a compression approach. In *International Conference on Machine Learning*, pages 254–263. PMLR.

Matan Ben Noach and Yoav Goldberg. 2020. [Compressing pre-trained language models by matrix de-](#)

[composition](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 884–889, Suzhou, China. Association for Computational Linguistics.

- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#).
- 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.
- Ondřej Dušek, Jekaterina Novikova, and Verena Rieser. 2020. [Evaluating the State-of-the-Art of End-to-End Natural Language Generation: The E2E NLG Challenge](#). *Computer Speech & Language*, 59:123–156.
- Jonathan Frankle and Michael Carbin. 2018. [The lottery ticket hypothesis: Training pruned neural networks](#). *CoRR*, abs/1803.03635.
- Mark Kurtz, Justin Kopinsky, Rati Gelashvili, Alexander Matveev, John Carr, Michael Goin, William Leiserson, Sage Moore, Bill Nell, Nir Shavit, and Dan Alistarh. 2020. [Inducing and exploiting activation sparsity for fast inference on deep neural networks](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 5533–5543, Virtual. PMLR.
- 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 *ICLR*. OpenReview.net.
- Namhoon Lee, Thalaiyasingam Ajanthan, and Philip Torr. 2018. [Snip: Single-shot network pruning based on connection sensitivity](#). In *International Conference on Learning Representations*.
- Yuchao Li, Fuli Luo, Chuanqi Tan, Mengdi Wang, Songfang Huang, Shen Li, and Junjie Bai. 2022. [Parameter-efficient sparsity for large language models fine-tuning](#). In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 4223–4229. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Chen Liang, Simiao Zuo, Minshuo Chen, Haoming Jiang, Xiaodong Liu, Pengcheng He, Tuo Zhao, and

- Weizhu Chen. 2021. [Super tickets in pre-trained language models: From model compression to improving generalization](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6524–6538, Online. Association for Computational Linguistics.
- David Lopez-Paz, Léon Bottou, Bernhard Schölkopf, and Vladimir Vapnik. 2015. Unifying distillation and privileged information. *arXiv preprint arXiv:1511.03643*.
- Ilya Loshchilov and Frank Hutter. 2017. [Fixing weight decay regularization in adam](#). *CoRR*, abs/1711.05101.
- Christos Louizos, Max Welling, and Diederik P Kingma. 2018. Learning sparse neural networks through l_0 regularization. *arXiv preprint arXiv:1712.01312*.
- Seyed-Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, and Hassan Ghasemzadeh. 2019. [Improved knowledge distillation via teacher assistant: Bridging the gap between student and teacher](#). *CoRR*, abs/1902.03393.
- Dragomir Radev, Rui Zhang, Amrit Rau, Abhinand Sivaprasad, Chiachun Hsieh, Nazneen Fatema Rajani, Xiangru Tang, Aadit Vyas, Neha Verma, Pranav Krishna, Yangxiaokang Liu, Nadia Irwanto, Jessica Pan, Faiaz Rahman, Ahmad Zaidi, Murori Mutuma, Yasin Tarabar, Ankit Gupta, Tao Yu, Yi Chern Tan, Xi Victoria Lin, Caiming Xiong, and Richard Socher. 2020. Dart: Open-domain structured data record to text generation. *arXiv preprint arXiv:2007.02871*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2018. [Language models are unsupervised multitask learners](#).
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100, 000+ questions for machine comprehension of text](#). *CoRR*, abs/1606.05250.
- Alex Renda, Jonathan Frankle, and Michael Carbin. 2020. [Comparing rewinding and fine-tuning in neural network pruning](#). *CoRR*, abs/2003.02389.
- Victor Sanh, Thomas Wolf, and Alexander Rush. 2020. [Movement pruning: Adaptive sparsity by fine-tuning](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 20378–20389. Curran Associates, Inc.
- 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](#). *CoRR*, abs/1909.05840.
- Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. 2019. [Patient knowledge distillation for BERT model compression](#). *CoRR*, abs/1908.09355.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Vladimir Vapnik. 1998. *Statistical learning theory*. Wiley.
- 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 *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). *CoRR*, abs/1804.07461.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. [Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers](#). *CoRR*, abs/2002.10957.
- Qingru Zhang, Simiao Zuo, Chen Liang, Alexander Bukharin, Pengcheng He, Weizhu Chen, and Tuo Zhao. 2022. [Platon: Pruning large transformer models with upper confidence bound of weight importance](#). In *International Conference on Machine Learning*, pages 26809–26823. PMLR.
- Giulio Zhou and Gerasimos Lampouras. 2020. [WebNLG challenge 2020: Language agnostic delexicalisation for multilingual RDF-to-text generation](#). In *Proceedings of the 3rd International Workshop on Natural Language Generation from the Semantic Web (WebNLG+)*, pages 186–191, Dublin, Ireland (Virtual). Association for Computational Linguistics.
- 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*.

A Results with More PLMs on subset of GLUE

In addition the widely used BERT and GPT-2 models, we also perform pruning experiments upon other two pre-trained language models: Electra_{base} and MiniLM_{12L-384H} to further verify the effectiveness of our method.

Due to computing resource constraint, we restrict our experiments on a subset of GLUE task, including RTE, CoLA and QNLI at 80% and 90% sparsity. We compare PINS against IMP and PLATON as two representative baselines for magnitude-based and sensitivity-based pruning methods. We fix the batch size as 32 and learning rate as $3e-5$ similar to the BERT experiments. We illustrate the pruning results on Table 8 and Table 9. At both sparsity levels, PINS consistently outperforms IMP and PLATON on all three datasets, verifying the general effectiveness of PINS for language model pruning.

Sparsity	Method	RTE Acc	CoLA Mcc	QNLI Acc
0%	Fine-tune	73.0	58.5	91.5
80%	IMP	60.5	21.6	87.5
	PLATON	68.2	54.1	89.8
	PINS (ours)	69.5	54.4	90.4
90%	IMP	57.5	14.1	83.9
	PLATON	63.1	38.8	88.0
	PINS (ours)	66.2	44.8	88.6

Table 8: Results with MiniLM_{12L-384H} on the GLUE development set.

Sparsity	Method	RTE Acc	CoLA Mcc	QNLI Acc
0%	Fine-tune	81.9	69.0	93.1
80%	IMP	59.9	11.2	87.5
	PLATON	73.6	60.0	91.0
	PINS (ours)	75.5	63.7	92.0
90%	IMP	52.9	0.0	83.0
	PLATON	69.9	48.0	89.7
	PINS (ours)	72.3	49.2	90.2

Table 9: Results with Electra_{base} on the GLUE development set.

B Proof of Theorem 1

Proof. Let t_i and t_j where $t_i \geq t_j$ denote the time steps at which two different checkpoints are saved; Let $R(f_{\theta^{(t \leftarrow t_i)}})$ and $R(f_{\theta^{(t \leftarrow t_j)}})$ denote the expected generalization error of models learned

from $f_{\theta^{(t_i)}}$ and $f_{\theta^{(t_j)}}$; Let n denotes the size of training data; $|\cdot|_{\mathcal{C}}$ denotes a capacity measure like VC-dimension for function class \mathcal{F}_{θ} . Based on previous expositions on VC theory, the following asymptotic generalization bound holds:

$$\begin{aligned}
 R(f_{\theta^{(t \leftarrow t_i)}}) &= R(f_{\theta^{(t \leftarrow t_i)}}) - R(f_{\theta^{(t_i)}}) \\
 &\quad + R(f_{\theta^{(t_i)}}) \\
 &\leq O\left(\frac{|\mathcal{F}_{\theta^{(t)}}|_{\mathcal{C}}}{n^{\alpha_i}}\right) + \epsilon_{t,t_i} + R(f_{\theta^{(t_i)}}) \\
 &= O\left(\frac{|\mathcal{F}_{\theta^{(t)}}|_{\mathcal{C}}}{n^{\alpha_i}}\right) + \underbrace{\inf_{f_{\theta^{(t)}} \in \mathcal{F}_{\theta^{(t \leftarrow t_i)}}} R(f_{\theta^{(t)}})}_{\text{bound}(f_{\theta^{(t \leftarrow t_i)}})} \\
 R(f_{\theta^{(t \leftarrow t_j)}}) &= R(f_{\theta^{(t \leftarrow t_j)}}) - R(f_{\theta^{(t_j)}}) \\
 &\quad + R(f_{\theta^{(t_j)}}) \\
 &\leq O\left(\frac{|\mathcal{F}_{\theta^{(t)}}|_{\mathcal{C}}}{n^{\alpha_j}}\right) + \epsilon_{t,t_j} + R(f_{\theta^{(t_j)}}) \\
 &= O\left(\frac{|\mathcal{F}_{\theta^{(t)}}|_{\mathcal{C}}}{n^{\alpha_j}}\right) + \underbrace{\inf_{f_{\theta^{(t)}} \in \mathcal{F}_{\theta^{(t \leftarrow t_j)}}} R(f_{\theta^{(t)}})}_{\text{bound}(f_{\theta^{(t \leftarrow t_j)}})}
 \end{aligned}$$

where ϵ_{t,t_i} is the approximation error of function class $\mathcal{F}_{\theta^{(t \leftarrow t_i)}}$ with respect to $f_{\theta^{(t_i)}}$. ϵ_{t,t_j} is defined in analogy. Because: (1) $\theta^{(t_i)}$ is a later checkpoint with higher sparsity than $\theta^{(t_j)}$, we have the learning speed $1 \geq \alpha_i \geq \alpha_j \geq \frac{1}{2}$; (2) $f_{\theta^{(t_i)}}$ has lower generalization error than $f_{\theta^{(t_j)}}$, we have the following inequality holds with high probability:

$$\text{bound}(f_{\theta^{(t \leftarrow t_i)}}) \leq \text{bound}(f_{\theta^{(t \leftarrow t_j)}})$$

□

C More Post-pruning Analyses

This section presents more visualized analyses of models sparsified by different pruning methods.

Figure 5 shows the layerwise rank distribution of BERT_{base} pruned using different importance criteria on the RTE dataset. The observation here is similar to what is discussed in the main body of the paper: PINS exhibits the lowest average matrix rank in the sparsified model compared to the other two criteria.

Figure 4 illustrates the weight distribution of BERT_{base} pruning using different importance criteria. From the left figure we can see that magnitude-based pruning tends to keep parameters with high absolute values, which is expected

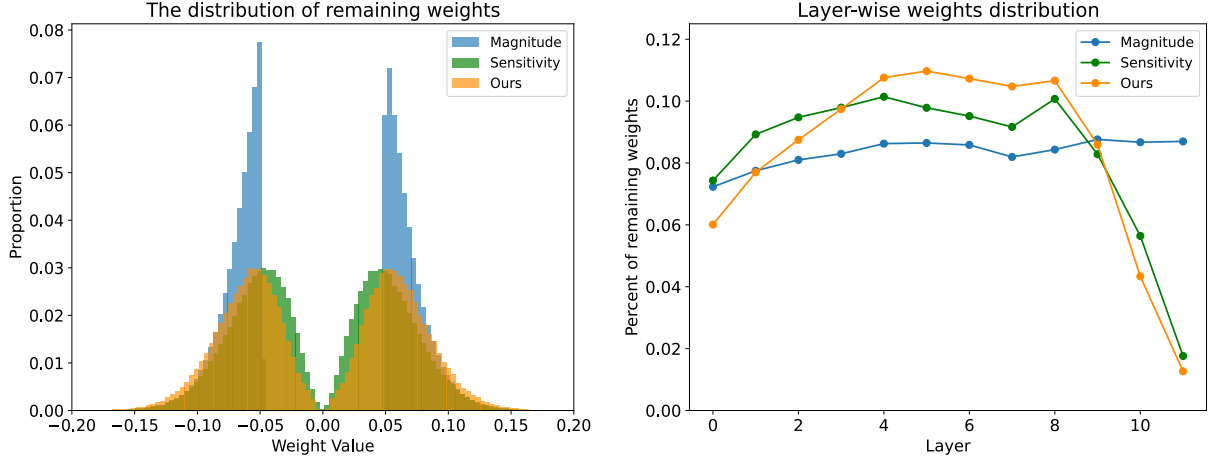


Figure 4: Weight distributions of $BERT_{\text{base}}$ pruned using different importance criteria on RTE dataset. Left figure shows the value distribution and the right figure shows how remaining parameters are distributed at different model layers.

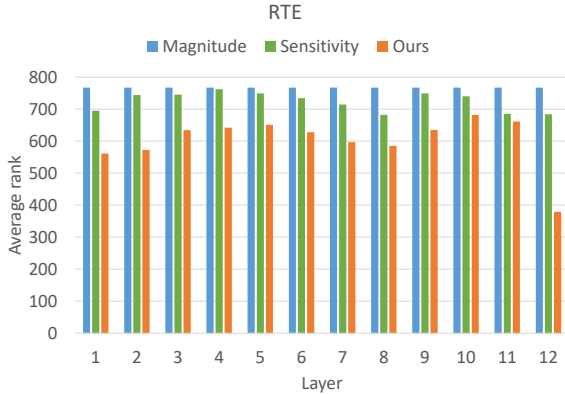


Figure 5: Layerwise rank distribution of $BERT_{\text{base}}$ pruning using different importance criteria on RTE dataset.

based on its definition. Sensitivity and PINS produce similar weight value distribution mainly because the two methods both contain the $g\theta$ term in their importance calculation. Despite the similarity, we can still observe that PINS produces smoother distribution than sensitivity and covers more weights with larger absolute values.

The right figure shows the layerwise distribution of remaining parameters after pruning. A clear trend is that PINS tends to retain more parameters in the middle layers (4-7), which also coincided with the inter-model sparsity pattern analysis in the main body of our paper. Both sensitivity and PINS remove a large proportion of parameters in the top layers (10-12) while magnitude-based pruning has no preference for model layers.

D Sparsity Scheduler

The proportion of remaining weights is controlled by the sparsity scheduler, here we adopt the commonly used cubic sparsity schedule to progressively reach target sparsity, i.e., $r^{(t)}$ at time step t within the maximum time steps T is given by:

$$\begin{cases} r_i & t \in [0, t_i) \\ r_f + (r_i - r_f) \left(\frac{T - t_f - t}{T - t_f - t_i} \right)^3 & t \in [t_i, T - t_f) \\ r_f & \text{otherwise} \end{cases} \quad (14)$$

where $r_i = 1.0$, r_f is the final percent of remained parameters, t_i and t_f are the warmup and cool-down steps.

E Accelerating Inference and Reducing Storage

We attain practical efficiency gain in terms of inference time and disk storage space using different sets of off-the-shelf techniques. Specifically, we use DeepSparse², a sparsity-aware inference runtime to accelerate inference of sparse model on CPUs. We also utilize the Pytorch built-in quantization function³ and Compressed Sparse Row (CSR) format⁴ to achieve a much smaller disk space requirement.

²<https://github.com/neuralmagic/deepsparse>

³<https://pytorch.org/docs/stable/quantization.html>

⁴https://github.com/huggingface/block-movement-pruning/blob/master/Saving_PruneBERT.ipynb

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
6
- A2. Did you discuss any potential risks of your work?
Not applicable. Left blank.
- A3. Do the abstract and introduction summarize the paper's main claims?
1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

4.1.1

- B1. Did you cite the creators of artifacts you used?
4.1.1
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Not applicable. Left blank.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Left blank.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
4.1.1;4.1.3

C Did you run computational experiments?

4

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
4.1.2

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

4.1.3

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

4.2

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

4.2

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Not applicable. Left blank.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Not applicable. Left blank.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Not applicable. Left blank.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Not applicable. Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Not applicable. Left blank.