

Pruning Large Language Models with Semi-Structural Adaptive Sparse Training

Weiyu Huang, Yuezhou Hu, Guohao Jian, Jun Zhu, Jianfei Chen[†]

Dept. of Comp. Sci. and Tech., Institute for AI, BNRist Center, THBI Lab,
Tsinghua-Bosch Joint ML Center, Tsinghua University
{hwy23, huyz21, jgh22}@mails.tsinghua.edu.cn; {dcszj, jianfeic}@tsinghua.edu.cn

Abstract

The remarkable success of Large Language Models (LLMs) relies heavily on their substantial scale, which poses significant challenges during model deployment in terms of latency and memory consumption. Recently, numerous studies have attempted to compress LLMs using one-shot pruning methods. However, these methods often suffer from considerable performance degradation on complex language understanding tasks, raising concerns about the feasibility of pruning in LLMs. To address this issue, we propose Adaptive Sparse Trainer (AST), a novel and efficient retraining framework tailored for semi-structured sparse models. AST enables models to learn optimal masks during the weight update process without incurring additional computational overhead. Furthermore, we demonstrate that incorporating knowledge distillation significantly improves retraining efficiency and enhances model performance under fixed computational constraints. Additionally, a supplementary set of well-initialized parameters is integrated to further augment the model’s efficacy. AST achieves state-of-the-art performance with minimal training cost. When applied to the LLaMA2-7B model, AST reduces the perplexity and zero-shot accuracy gap between dense and 2:4 semi-structured sparse models to 0.6 and 1.16%, respectively, utilizing less than 0.4% of the pretraining tokens and GPU hours. Our work demonstrates the feasibility of deploying semi-structured sparse LLMs and offers a promising alternative for achieving highly compressed models when combined with existing quantization techniques.

Code — <https://github.com/thu-ml/Adaptive-Sparse-Trainer>

Extended Version — <https://arxiv.org/abs/2407.20584>

1 Introduction

Transformer-based Large Language Models (LLMs) are equipped to handle complex tasks (Devlin et al. 2018; Brown et al. 2020; Achiam et al. 2023) and exhibit emergent abilities (Wei et al. 2022) due to their expanding parameter count. However, this continual growth in model size poses significant challenges for practical deployment. In particular, inference speed suffers due to the increasing computational and memory demands. This has spurred a series of efforts to develop effective model compression techniques

aimed at reducing memory footprints and easing the constraints associated with deploying these large-scale models.

Model pruning (Frantar and Alistarh 2023; Han, Mao, and Dally 2016; Sun et al. 2023) is an effective technique for compressing LLMs by setting a proportion of weights to zero, thereby adhering to a specific sparsity pattern. Recently, N:M sparsity has emerged as a type of semi-structured sparsity pattern that offers an optimal balance between precision and hardware efficiency. Specifically, N:M sparsity retains only N nonzero elements out of every group of M elements. This sparsity pattern can accelerate both matrix multiplication and memory access, potentially enhancing the performance of both pre-filling and decoding processes on off-the-shelf GPUs. Despite the promise of N:M sparsity, current state-of-the-art methods for pruning LLMs, such as SparseGPT (Frantar and Alistarh 2023) and Wanda (Sun et al. 2023), employ a post-training approach that determines the sparsity pattern in a layer-by-layer fashion without back-propagation. Although these methods improve efficiency, they can lead to significant performance degradation, particularly in knowledge-intensive tasks (Nowak et al. 2024), raising concerns about the feasibility of pruning LLMs. Additionally, while retraining pruned models has been successful in the pre-LLM era (Wang, Wohlwend, and Lei 2019; Lee, Ajanthan, and Torr 2019; Evci et al. 2020), its application to models with billions of parameters remains under-explored.

Although retraining sparse LLMs holds significant potential, it introduces several unique challenges: (1) Retraining is computationally expensive, necessitating techniques that ensure rapid convergence. (2) The output model must adhere to strict sparsity pattern, which adds additional constraints to the retraining process. (3) In order to achieve optimal performance, both masks and weights should be learnable during training. (4) Pruning LLMs may compromise essential language understanding and reasoning abilities, which are significantly harder to restore compared to simpler metrics such as perplexity. To this end, we propose a novel and efficient training method, Adaptive Sparse Trainer (AST), designed to produce high-performance sparse LLMs. *For the first time, we demonstrate that pretrained 2:4 sparse LLMs can achieve competitive performance not only in terms of perplexity but also in more demanding knowledge-intensive tasks, making them viable for practical deployment.* AST

[†]Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

integrates the transition from dense to sparse models directly into the training process, gradually decaying unimportant weights to zero through a carefully designed decay mechanism. This approach allows for the revival of pruned weights during training, enabling a smooth and dynamic adjustment of the model’s connectivity pattern while adhering to the N:M sparsity structure. Furthermore, AST applies knowledge distillation using dense model as the teacher, which can significantly speed up model convergence and prevent sparse models from settling into local optima. It also helps the sparse model retain the valuable world knowledge and performance characteristics of the original dense model, thereby enhancing its generalization ability and compensating for the use of weaker training datasets. To further enhance performance, a supplementary set of parameters is integrated using information from pruned weights. When applied to the LLaMA2-7B model, AST achieves a 2:4 sparse configuration with minimal performance loss, demonstrating that compressed models can still perform effectively in practice. Our model further benefits from AWQ quantization, achieving competitive compression rates with state-of-the-art performance.

The key contributions of this work are as follows:

- We propose Adaptive Sparse Trainer (AST), an novel semi-structured pruning framework designed to compress large language models efficiently while maintaining high performance.
- AST introduces a gradual decay scheduler (Annealing SR-STE) combined with knowledge distillation to accelerate model convergence and improve performance on complex tasks.
- We introduce SLoRB, a technique that boosts model performance by adding a supplementary set of well-initialized parameters.
- When applied to LLaMA2-7B model, our 2:4 sparse model experiences only a 0.6 increase in perplexity and a 1.16% accuracy drop in zero-shot tasks, with less than 0.4% of the pretraining cost.

2 Related Work

Network Pruning Pruning is a prevalent technique aimed at reducing model size and computational costs. It originated from methods like OBD (LeCun, Denker, and Solla 1990) and OBS (Hassibi, Stork, and Wolff 1993). Based on sparsity patterns, pruning methods can be broadly categorized into unstructured, structured, and semi-structured pruning. Unstructured pruning removes individual weights (Han, Mao, and Dally 2016; Paul et al. 2023), which can maintain performance even with high sparsity. However, due to its random pattern, unstructured models are difficult to accelerate. Structured pruning (Liu et al. 2017; Molchanov et al. 2019; Nova, Dai, and Schuurmans 2023; Shen et al. 2022), on the other hand, removes entire neurons, filters, or attention heads, resulting in models that are easier to accelerate on standard hardware but often suffer from severe performance loss. Semi-structured sparsity (e.g., N:M sparsity) (Hubara et al. 2021) has been applied as a trade-off between

performance and achieving actual speedup. Recently, a series of works (Frantar and Alistarh 2023; Sun et al. 2023; Zhang et al. 2024a,b) have made progress in pruning LLMs with billions of parameters. However, these pruned models from training-free methods still fall short in complex zero-shot performance.

Retraining Pruned Models Another line of research (Han et al. 2015; Singh and Alistarh 2020; Renda, Frankle, and Carbin 2020; Zhou et al. 2023) has focused on retraining pruned models to enhance performance. While retraining has been shown to be effective with smaller models and simpler tasks (Kurtic and Alistarh 2022; Zhu and Gupta 2018), it often involves additional training steps (Frankle and Carbin 2019) or introduces extra parameters for the pruning process (Shi et al. 2023), which limits its application to large-scale models due to the substantial computational resources required. Other retraining methods focus on unstructured sparsity, which may struggle when given stricter sparsity patterns (Lee, Ajanthan, and Torr 2019; Evci et al. 2020). Recently, Sheared LLaMA (Xia et al. 2024) employed a two-stage structured pruning process to train models that outperform others of similar size. However, they employ a computationally intensive pruning stage, unsuitable for more fine-grained sparsity. Our work proposes a lightweight pruning process that enables the retraining of semi-structured sparse large language models with minimal training costs.

Combining Pruning with Quantization Pruned models can be further compressed using quantization. Earlier methods like Deep Compression (Han, Mao, and Dally 2016) combined pruning and quantization to significantly reduce the size of neural networks. More recent work (Frantar and Alistarh 2023; Sun et al. 2023; Guo et al. 2024) has combined sparsity with quantization in large language models. In our work, we report results using AWQ quantization (Lin et al. 2024) with our semi-structured sparse model.

3 Methods

We begin by revisiting a naive approach to training sparse models before introducing our method. The full training pipeline is illustrated in Figure 1.

3.1 Sparse Training

Pruning model weights is equivalent to multiplying them by an element-wise mask. For example, consider the matrix multiplication in a linear layer:

$$Z = XW^T, \quad Z \in \mathbb{R}^{N \times D}, \quad X \in \mathbb{R}^{N \times C}, \quad W \in \mathbb{R}^{D \times C},$$

where X , W , and Z represent the model input, weight matrix, and output activation, respectively. The pruned weight matrix can be expressed as:

$$\tilde{W} = m(W) \odot W, \quad m(W) \in \{0, 1\}^{D \times C},$$

where $m(\cdot)$ is a mapping that selects a mask based on the weight matrix W . In this work, we focus on N:M sparsity (Hubara et al. 2021), where there are N nonzero elements in every M consecutive weights in the same row. However, when implementing the backward pass for a sparse

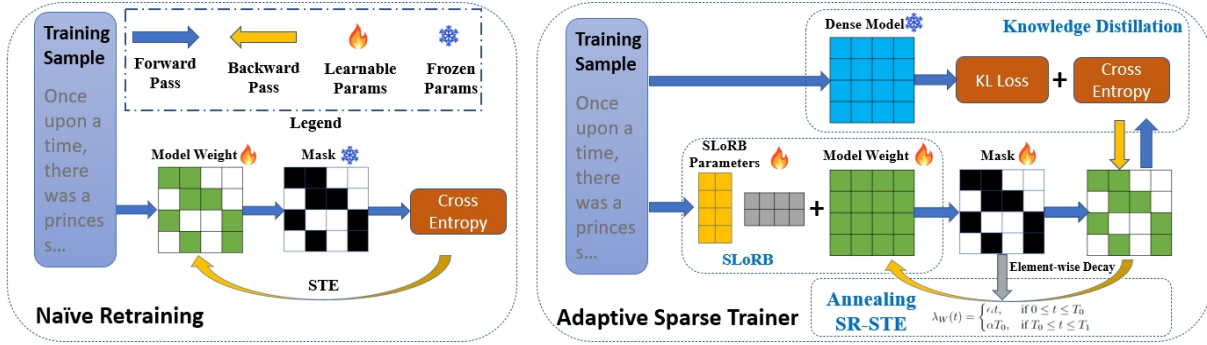


Figure 1: **(Left)** In the naive training baseline, the mask remains constant during training, which can result in suboptimal performance. **(Right)** Adaptive sparse training enables both mask and weight learning through a scheduled decay term. AST also utilizes distillation and SLoRB parameters to speed up convergence and improve performance.

model with automatic differentiation, the gradient cannot flow through the discrete mask $m(W)$. A straightforward approach to solve this issue is the straight-through estimator (STE) (Bengio, Léonard, and Courville 2013) which updates the parameters using the gradient with respect to the masked weight \tilde{W}_t , where W_t is the parameter at iteration t :

$$W_{t+1} \leftarrow W_t - \gamma_t g(\tilde{W}_t), \quad (1)$$

here $g(\tilde{W}_t)$ represents the gradient with respect to \tilde{W}_t and γ_t indicates the learning rate at iteration t .

A previous method (Sun et al. 2023) employs a fixed mask derived from one-shot pruning techniques and updates only the remaining parameters using language modeling loss. However, this strategy has two significant limitations. First, it discards the valuable information embedded in the pruned weights and prevents a smooth transition from a dense to a sparse model, resulting in slower convergence and suboptimal performance. Second, relying solely on the language modeling loss during retraining increases the risk of the sparse model becoming trapped in a local optimum, leading to poor outcomes.

To address these issues, we maintain all model weights and select the mask on-the-fly while gradually decaying non-essential weights to zero, thereby enabling a smooth transition (§3.2). Additionally, we leverage knowledge distillation to guide the model toward a global optimum, significantly accelerating convergence under a fixed computational budget (§3.3). Finally, we offer an optional method to further enhance model performance by incorporating an additional set of well-initialized parameters, as detailed in (§3.4).

3.2 Adaptive Mask Learning with Annealing SR-STE

As highlighted in previous work (Frankle and Carbin 2019; Evcı et al. 2020), the connectivity pattern is as crucial as the weights themselves in determining the performance of sparse models. Therefore, we enable the model to adaptively learn the optimal mask during training instead of fixing it permanently. Specifically, we recalculate $m(W_t)$ based on the magnitude criterion every Δt iterations and add a decay term to the masked weights. The intuition behind this

approach is straightforward: important weights typically receive strong gradients of the same sign during training, causing their magnitude to grow despite the decay, and they will eventually be revived. In contrast, less important weights tend to receive weaker, mixed gradient signals, leading them to decay to zero and remain masked under our settings. We build upon SR-STE (Zhou et al. 2021) by applying L2 decay to masked weights, as this approach helps preserve overall model performance.

However, selecting the appropriate decay strength is challenging. If the decay signal is too strong, masked weights will struggle to regenerate, leading to an almost fixed mask. On the other hand, if the decay signal is too weak, the model will fail to filter out important weights, resulting in heavy oscillation that hinders model convergence. To address this dilemma between stability and exploration, we propose adding a moving decay schedule, with weaker decay strength at the beginning to encourage the model to explore various connectivity patterns and a stronger decay signal toward the end to facilitate model convergence.

Formally, we update the model weights with the following equation:

$$\lambda_W(t) = \begin{cases} \alpha t, & \text{if } 0 \leq t \leq T_0, \\ \alpha T_0, & \text{if } T_0 \leq t \leq T_1, \end{cases} \quad (2)$$

$$W_{t+1} \leftarrow W_t - \gamma_t \left(g(\tilde{W}_t) + \lambda_W(t) \left(\overline{m(W_t)} \odot W_t \right) \right), \quad (3)$$

where $\lambda_W(t)$ denotes the decay factor, $\overline{m(W_t)}$ denotes the mask for pruned weights at iteration t , T_1 denotes the total number of training batches, and T_0 represents the batch at which the decay factor stops increasing. Compared with STE in Equation 1, masked weights receive a decay signal proportional to the decay factor $\lambda_W(t)$, which starts small at the beginning and increases later on. This method, termed Annealing SR-STE (ARS-STE), has been shown to improve model performance compared to naive SR-STE. We also provide a detailed analysis of mask oscillation behavior during training in Section 7 of the Appendix.

We experimented with various one-shot pruning criteria incorporating activation information (Frantar and Alistarh 2023; Zhang et al. 2024a) for mask selection. However, the *classic magnitude criterion consistently outperformed these methods in both computational efficiency and performance*. As activation-based approaches incur higher computational costs and are less reliable due to activation fluctuations caused by weight updates, leading to degraded performance in later training stages.

Since the naive SR-STE is only compatible with the SGD optimizer, we introduce further adjustments to support more advanced optimizers such as AdamW. Specifically, as AdamW maintains both first-order and second-order moments of the gradient, we decouple the weight decay term from the first-order moment instead of directly adding decay to the gradient, as done in (Hu et al. 2024). This separation ensures that the weight decay term depends solely on the current weight, avoiding interference with momentum calculations. We then use the decoupled first-order signal to update the second-order moment. Detailed expressions and explanations are provided in Section 6 of the Appendix.

3.3 Alleviating the Retraining Dilemma through Knowledge Distillation

A key distinction between pretraining and retraining pruned models lies in how their parameters are initialized. In the pretraining setting, parameters are typically randomly sampled, whereas in retraining, they are inherited from a well-trained model. Consequently, pruned models retain a portion of the patterns acquired during pretraining.

We discovered that this feature tends to trap pruned models in local optima. Specifically, while retraining can achieve much faster initial convergence, it often fails to reach an optimal state later on; we refer to this phenomenon as the **Retraining Dilemma**. As illustrated in Figure 2, although using cross-entropy loss during retraining initially leads to a quicker reduction in training loss compared to pretraining, the test set perplexity remains unstable and elevated. We attribute this to the fact that, unlike randomly initialized models, the pruned model is more prone to overfitting on the limited data available at the start due to its prior knowledge. This overfitting hampers the model’s ability to learn global features and achieve optimal convergence. Notably, this issue persists even when using a significantly smaller learning rate than in pretraining, suggesting that simply reducing the learning rate is insufficient to resolve the problem.

As a solution, we found that applying KL-divergence loss (Kullback and Leibler 1951) can address this issue. Unlike the language modeling loss, KL-divergence loss measures the difference between two probability distributions, offering richer signals that help mitigate overfitting in the early stages of training. Consequently, we employ the following loss function:

$$\mathcal{L}_{logit} = D_{KL}(p_{\theta_t} || p_{\theta_s}) = \frac{1}{B \times S} \sum_{i=1}^{B \times seq} p_{\theta_t}(x_i) \log \frac{p_{\theta_t}(x_i)}{p_{\theta_s}(x_i)},$$

$$\mathcal{L} = \alpha \mathcal{L}_{logit} + (1 - \alpha) \mathcal{L}_{task}, \quad (4)$$

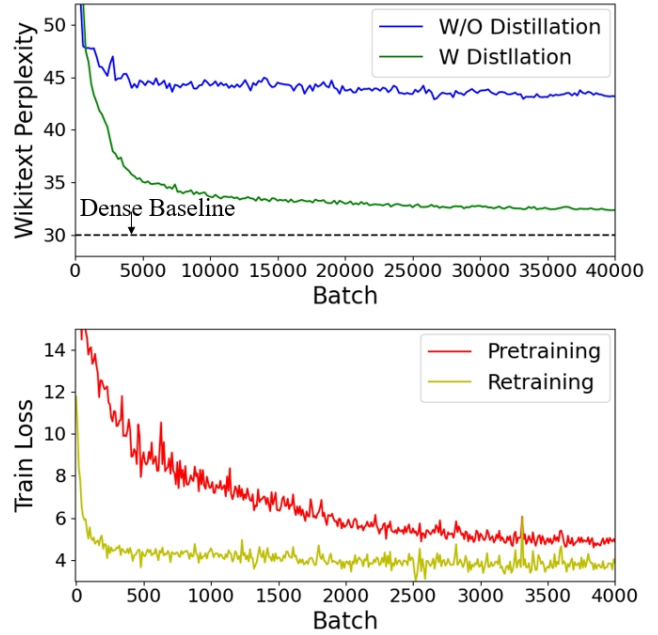


Figure 2: (Upper) The Wikitext perplexity curve for retraining GPT2 with and without knowledge distillation. (Lower) The training loss curve for pretraining and retraining.

where \mathcal{L}_{task} is the cross-entropy loss, B is the batch size, S is the sequence length, and p_{θ_t} and p_{θ_s} are the probability distributions of the teacher and student models, respectively.

We observe that the Retraining Dilemma is more pronounced in smaller models; therefore, we typically apply a larger α for these models. We also experimented with various distillation methods that utilize intermediate information for distillation. However, we found that adding constraints to intermediate outputs hinders the model’s generalization ability and leads to undesirable outcomes.

3.4 Sparse Low-Rank Boosting

The pruned model with 2:4 sparsity, which retains only half of the parameters under strict structural constraints, experiences a reduction in expressive capacity. To mitigate this, we incorporate a LoRA (Low-Rank Adaptation) (Hu et al. 2022) shape parameter that is trained alongside the original parameters, helping to bridge the performance gap between dense and sparse models with minimal memory overhead. Unlike traditional LoRA fine-tuning, where the original model parameters are frozen and only the adapters are trained, our approach trains both the sparse parameters and the adapter weights simultaneously. This approach recognizes that while the classic LoRA method freezes the original model to prevent overfitting and catastrophic forgetting during downstream tasks, our method enhances generalization by training both the additional and original weights on a pretraining dataset, thereby enabling concurrent training.

Another notable aspect of our method is the initialization strategy. Classic LoRA typically employs random initialization techniques, such as Xavier initialization (Glorot




Figure 3: **Visualization of SLoRB Initialization:** Consider a weight matrix of size 3 by 8, with $k = 4$. Weight S_{ij} is broadcasted within group G_{ij} after multiplication with the projection matrix.

and Bengio 2010) or Kaiming initialization (He et al. 2015). However, in our approach, we leverage the pruned weights as additional information to initialize the adapter weights, thereby accelerating the training process. Given that in 2:4 sparsity, each neuron can only access half of the inputs, we find that incorporating additional information to retain the weight’s first-order information helps preserve the model’s capacity. Specifically, for a weight matrix W of size $N \times d$ and its corresponding mask matrix M , we select the rank r as $\frac{d}{k}$, where k can be 64, 32, 16, and so on. A smaller k improves performance but also increases memory usage. We select a projection matrix X of size $\frac{d}{k} \times d$ and a SLoRB weight matrix S of size $N \times \frac{d}{k}$. The matrices X and S are defined as follows:

$$x_{ij} = \begin{cases} 1, & \text{if } i \cdot k \leq j \leq (i + 1) \cdot k - 1, \\ 0, & \text{otherwise,} \end{cases} \quad (5)$$

$$S_{ij} = \frac{1}{k} \sum_{p=j \cdot k}^{(j+1) \cdot k - 1} W_{ip} \cdot \neg M_{ip}. \quad (6)$$

We define Group G_{ij} as the elements from $j \cdot k$ to $(i + 1) \cdot k - 1$ in row i . As illustrated in Fig 3, each SLoRB weight S_{ij} is broadcast within Group G_{ij} . By setting S_{ij} as the mean of all pruned weights in Group G_{ij} , this design ensures that the total mean of groups G_{ij} remains consistent after pruning. Although different initialization methods may ultimately converge to similar outcomes given sufficient training steps, our method converges more rapidly, making it particularly advantageous when the computational budget is limited, as demonstrated in Section 4 of the Appendix. We refer to this method as **Sparse Low-Rank Boosting (SLoRB)**, which is an optional technique that trades off memory overhead for enhanced performance. The complete pseudo-code for our method is provided in Algorithm 1.

4 Experiments

4.1 Experiment Setup

Model Configurations. We report the performance of Adaptive Sparse Trainer (AST) across three different LLM model families: LLaMA2-7B (Touvron et al. 2023), OPT (Zhang et al. 2022), and GPT2 (Brown et al. 2020). We

Algorithm 1: Training Process for AST

Input: Weight of linear layers $\mathcal{W}^{(0)}$; Mask update frequency Δt ; Total training iteration T_0 ; SLoRB parameter k ; Decay increase iteration T_1 ;

- 1: **for** $W^{(0)} \in \mathcal{W}^{(0)}$ **do**
 - 2: Initialize the mask $m(W^{(0)})$ based on magnitude;
 - 3: **if** Use SLoRB **then**
 - 4: Initialize SLoRB weight $S^{(0)} \in \mathcal{S}^{(0)}$ and $X^{(0)} \in \mathcal{X}^{(0)}$ by Equation 5 and 6;
 - 5: **end if**
 - 6: **end for**
 - 7: **for** $t = 1, 2, \dots, T_0$ **do**
 - 8: **if** $t \bmod \Delta t == 0$ **then**
 - 9: **for** $W^{(t)} \in \mathcal{W}^{(t)}$ **do**
 - 10: Update model mask $m(W^{(t)})$;
 - 11: **end for**
 - 12: **end if**
 - 13: Compute decay for term $\lambda_W(t)$ from Equation 2;
 - 14: Compute gradient $g(W^{(t)})$ by back-propagation on distillation loss \mathcal{L} from Equation 4;
 - 15: Update weight $W^{(t)}$ by Equation 3;
 - 16: **if** Use SLoRB **then**
 - 17: Update $S^{(t)}$ and $X^{(t)}$ by gradient;
 - 18: **end if**
 - 19: **end for**
 - 20: **return** the pruned model.
-

present results for two different sparsity patterns: AST-Naive, which adheres to a strict 2:4 sparsity pattern without additional weights, and AST-Boosted, which incorporates extra SLoRB weights. For AST-Boosted models, we selected $k = 16$, introducing an additional 12.5% of parameters. Optimal hyperparameters were identified through a grid search, with the specific hyperparameters and training details provided in Section 3 of the Appendix.

Data. For training smaller models like the OPT and GPT2 model families, we utilized the C4 dataset (Raffel et al. 2020). For the LLaMA2-7B model, we employed a more comprehensive dataset, RedPajama-v1¹, which encompasses data from seven domains: CommonCrawl, C4, GitHub, Wikipedia, Books, ArXiv, and StackExchange. Additionally, we leveraged the dynamic batching feature provided in the ShearedLLaMA (Xia et al. 2024) codebase.

Baseline. For the GPT2 and OPT models, we compare our methods with both training-free and training-based approaches. Among training-free methods, we include comparisons with SparseGPT (Frantar and Alistarh 2023) and Wanda (Sun et al. 2023). For training-based methods, given the lack of existing results for retraining generative language models with N:M sparsity, we implemented iterative magnitude pruning (Frankle and Carbin 2019) and gradual magnitude pruning (Kurtic and Alistarh 2022), both of which can be adapted to achieve the desired sparsity. To ensure a fair comparison in terms of computational cost, we re-

¹<https://huggingface.co/datasets/togethercomputer/RedPajama-Data-1T>

Method	Training	OPT			GPT2			
		125M	350M	1.3B	124M	350M	774M	1.5B
Dense	-	27.76	22.00	14.62	29.95	21.72	19.43	17.40
SparseGPT	✗	45.58	40.33	29.03	50.09	31.03	25.98	21.14
Wanda	✗	60.91	50.16	23.92	115.64	63.71	49.97	30.44
Iterative Magnitude Pruning	✓	38.37	30.29	23.94	40.08	29.86	24.31	20.83
Gradual Magnitude Pruning	✓	31.51	25.98	16.78	33.48	24.83	22.01	18.96
AST-Naive(Ours)	✓	30.22	24.65	15.85	32.23	23.65	21.29	18.33
AST-Boosted(Ours)	✓	28.68	24.03	15.43	31.13	23.03	20.66	18.01

Table 1: Perplexity results on raw-WikiText2 on 2:4 sparsified language models. AST outperforms both training-free and other training-based methods with similar computational costs.

Models (Sparsity Pattern)	Parameters	BoolQ	RTE	HellaSwag	WinoGrande	ARC-e	ARC-c	OBQA	Mean
LLaMA2-7B (Dense)	6.7B	77.73	63.89	57.18	69.04	76.17	42.91	31.60	59.78
Sheared-LLaMA-2.7B (Dense)	2.7B	65.99	50.54	51.21	65.03	67.29	33.27	28.80	51.73
INCITE-Base-3B (Dense)	2.8B	67.40	52.34	47.91	62.98	67.68	31.74	27.60	51.09
Open-LLaMA-3B-v2 (Dense)	3.4B	65.89	55.69	52.12	63.59	68.34	34.32	26.00	52.13
LLaMA-7B (Sparse)	3.4B	73.21	61.34	54.86	66.18	70.24	35.68	31.80	56.19
LLaMA2-7B (Sparse)	3.4B	73.12	66.06	54.66	67.87	73.61	39.93	28.60	57.68
LLaMA2-7B (Sparse+SLoRB)	4.2B	75.04	66.06	55.24	68.48	74.91	41.11	29.40	58.62

Table 2: Accuracy (%) of various open-sourced models on seven zero-shot tasks.

port results for training-based methods that include distillation, as we have observed that incorporating distillation significantly enhances performance. Due to computational constraints, we report the results of training-based baselines only for the GPT2 and OPT models.

For the LLaMA2-7B model, we compare our approach against strong dense pre-trained LLMs with a similar number of non-zero parameters. Additionally, we include a comparison with the LLaMA-7B 2:4 sparsity model provided in the Wanda (Sun et al. 2023) paper.

Evaluation. We evaluated the WikiText-2 perplexity for all models and assessed zero-shot and few-shot performance on LLaMA2-7B using EleutherAI’s LM Harness (Gao et al. 2021). Our evaluation tasks included zero-shot ARC Easy (Clark et al. 2018), OpenBookQA (Mihaylov et al. 2018), WinoGrande (Sakaguchi et al. 2021), RTE from the GLUE benchmark (Wang et al. 2018), HellaSwag (Zellers et al. 2019), ARC Challenge (Clark et al. 2018), BoolQ (Clark et al. 2019), as well as more complex tasks like the MMLU benchmark (Hendrycks et al. 2021a), GSM8K (Cobbe et al. 2021), and MATH (Hendrycks et al. 2021b).

4.2 Main Results

Language Modeling In Table 1, we compare the perplexity results of our sparse models with those of the baselines. Our methods significantly outperform both training-free and training-based approaches across various models, achieving results close to those of the dense baseline. A plausible reason why training-based methods like IMP fail is that while a lottery ticket may exist for a randomly initialized model, this may not be the case for a well-trained model. As for the

gradual baseline, applying a proportion of training steps at low sparsity may not provide the model with sufficient time to converge.

Zero-shot Result In Table 2, we present the accuracy results of seven zero-shot tasks for AST and baseline models. Our method performed best in most of the tasks. It should be noted that our sparse models consistently outperform smaller dense models with a similar parameter count, suggesting a promising approach for obtaining parameter-efficient models. Moreover, our method requires only a minimal amount of training tokens to converge. For example, when training the LLaMA2-7B model, **we utilized only 7B tokens, which is less than 0.4% of those used in pretraining, making AST applicable to open-source models.**

Our method can also be seamlessly adapted to current quantization techniques to achieve extremely compressed models with minimal performance drop. We provide results using AWQ quantization methods in section 8 of the Appendix.

4.3 Additional Results for LLaMA2-7B

Recent findings (Nowak et al. 2024) have shown that previous pruning methods suffer significant performance degradation in knowledge-intensive tasks. To address this, we tested our LLaMA2-7B model on more complex tasks. As shown in Table 3, our model preserves most of the knowledge from the dense model and continues to perform well in knowledge-intensive tasks compared with previous pruning methods. This provides strong evidence that 2:4 pruned models possess considerable potential, contrary to the observations in previous studies.

LLaMA2-7B	Dense	Wanda	AST-Naive	AST-Boosted
Perplexity ↓	5.12	11.02	5.82	5.69
MMLU (5-shot) ↑	45.3	27.6	37.9	38.2
MATH (4-shot) ↑	5.38	2.86	4.42	4.64
GMS8K (LoRA Finetuned) ↑	40.3	32.1	35.6	36.2

Table 3: Results of perplexity and knowledge-intensive tasks for LLaMA2-7B models with 2:4 sparsity. The symbols ↓ and ↑ indicate that lower and higher values are better, respectively. (LoRA methods are finetuned with $r=64$)

Method	GPT			OPT
	124M	350M	774M	125M
Dense	29.95	21.72	19.43	27.76
Naive Training	40.34	29.79	28.65	39.46
No distillation	39.29	29.08	27.21	36.97
Static SR-STE	32.84	24.04	21.73	31.08
Fixed Mask	32.93	24.18	21.95	31.06
AST-Naive(Ours)	32.23	23.65	21.29	30.22

Table 4: Ablation study of different methods in training sparse models.

4.4 Ablation Study

Training Sparse Models. Our AST-Naive method employed distillation, Annealing SR-STE, and adaptive masking during training. In Table 4, we present an ablation study to assess the effects of each of these components. For a fair comparison, we apply additional training steps to non-distillation methods. Specifically, we analyze the impact of training without distillation, using a naive SR-STE decay factor, and fixing the mask during training, each component in isolation. Additionally, we provide results for the naive training baseline in Figure 1 mentioned above. We demonstrate that all three approaches contribute to performance gains compared to naive training across various models.

Distillation Functions We also conducted ablation studies on different distillation functions used in previous work, including TinyBERT (Jiao et al. 2020), MobileBERT (Sun et al. 2020), and Sparse-Finetuning (Kurtic et al. 2023), which use attention and hidden states for distillation. Additionally, we examined MiniLLM (Gu et al. 2024), which employs reverse KL divergence. We find that using intermediate information during the distillation of generative language models is detrimental, and that KL loss is sufficient for optimal performance. Reverse-KL yields performance similar to that of forward KL. Detailed descriptions of the distillation loss functions for each of these methods are provided in Section 5 of the Appendix.

4.5 Speedup

Table 6 presents speedup results obtained using TensorRT-LLM². We evaluate the actual end-to-end decoding speedup on two GPU architectures with the LLaMA2-7B 2:4 sparse model. We employ throughput, measured as the number of

²<https://github.com/NVIDIA/TensorRT-LLM>

Method	GPT			OPT
	124M	350M	774M	125M
Dense	29.95	21.72	19.43	27.76
TinyBERT	42.75	33.35	30.86	39.39
MobileBERT	44.87	31.67	29.75	40.33
Sparse-Finetuning	41.19	29.42	26.19	38.96
MiniLLM	32.20	23.64	21.31	30.24
KL Loss(Ours)	32.23	23.65	21.29	30.22

Table 5: Ablation study on different distillation loss for training sparse models.

RTX4090			
Inp Len, Out Len	Sparse	Dense	Speedup
128, 128	70.23	52.94	1.33x
128, 1024	69.11	52.00	1.33x
1024, 128	68.06	51.10	1.33x
1024, 1024	67.41	50.37	1.34x
L20			
Inp Len, Out Len	Sparse	Dense	Speedup
128, 128	54.75	29.86	1.83x
128, 1024	53.81	29.57	1.82x
1024, 128	52.49	29.18	1.80x
1024, 1024	51.64	28.94	1.78x

Table 6: Speedup results using TensorRT-LLM on RTX4090 and L20 GPUs with different input and output sequence lengths, measured by throughput (tokens/s).

tokens processed per second, as the primary evaluation metric. Across a range of input and output lengths, the 2:4 sparse model demonstrates an overall acceleration of 1.33x to 1.83x compared to its dense counterpart, highlighting its potential for practical deployment.

5 Conclusion

In this paper, we introduce the Adaptive Sparse Trainer (AST), a novel and efficient training pipeline for semi-structured sparse models. AST effectively narrows the precision gap between dense and sparse LLMs in terms of perplexity and accuracy on zero-shot tasks, while keeping training costs minimal. Our results demonstrate that pruning LLMs is feasible with minimal performance loss on knowledge-intensive tasks, and that large semi-structured sparse models can outperform dense models of similar decoding speed when supported by the appropriate software and hardware. Although our findings contribute to advancing the retraining of pruned models with billions of parameters, we acknowledge that the limited number of training tokens, due to computational constraints, remains an area for future exploration. Expanding this work to larger models or increasing training tokens count could provide valuable insights and further enhance the effectiveness of our methods.

Acknowledgements

The authors would like to thank Ziteng Wang, Bingrui Li and Pengle Zhang for helpful discussions and suggestions on code implementation. This work was supported by NSFC Project (No. 62376131).

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Bengio, Y.; Léonard, N.; and Courville, A. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *NeurIPS*, 1877–1901.
- Clark, C.; Lee, K.; Chang, M.-W.; Kwiatkowski, T.; Collins, M.; and Toutanova, K. 2019. BoolQ: Exploring the surprising difficulty of natural Yes/No questions. In *ACL*, 2924–2936.
- Clark, P.; Cowhey, I.; Etzioni, O.; Khot, T.; Sabharwal, A.; Schoenick, C.; and Tafford, O. 2018. Think you have solved question answering? Try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Evcı, U.; Gale, T.; Menick, J.; Castro, P. S.; and Elsen, E. 2020. Rigging the lottery: Making all tickets winners. In *ICML*, 2943–2952.
- Frankle, J.; and Carbin, M. 2019. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *ICLR*.
- Frantar, E.; and Alistarh, D. 2023. Sparsegpt: Massive language models can be accurately pruned in one-shot. In *ICML*, 10323–10337.
- Gao, L.; Tow, J.; Biderman, S.; Black, S.; DiPofi, A.; Foster, C.; Golding, L.; Hsu, J.; McDonell, K.; Muennighoff, N.; et al. 2021. A framework for few-shot language model evaluation. *Version v0. 0.1. Sept*, 8.
- Glorot, X.; and Bengio, Y. 2010. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, 249–256.
- Gu, Y.; Dong, L.; Wei, F.; and Huang, M. 2024. MiniLLM: Knowledge distillation of large language models. In *ICLR*.
- Guo, J.; Wu, J.; Wang, Z.; Liu, J.; Yang, G.; Ding, Y.; Gong, R.; Qin, H.; and Liu, X. 2024. Compressing large language models by joint sparsification and quantization. In *ICML*.
- Han, S.; Mao, H.; and Dally, W. J. 2016. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *ICLR*.
- Han, S.; Pool, J.; Tran, J.; and Dally, W. J. 2015. Learning both weights and connections for efficient neural networks. In *NeurIPS*, 1135–1143.
- Hassibi, B.; Stork, D. G.; and Wolff, G. J. 1993. Optimal brain surgeon and general network pruning. In *ICNN*, 293–299.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, 1026–1034.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2021a. Measuring massive multitask language understanding. In *ICLR*.
- Hendrycks, D.; Burns, C.; Kadavath, S.; Arora, A.; Basart, S.; Tang, E.; Song, D.; and Steinhardt, J. 2021b. Measuring Mathematical Problem Solving With the MATH Dataset. *arXiv:2103.03874*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-rank adaptation of large language models. In *ICLR*.
- Hu, Y.; Zhao, K.; Huang, W.; Chen, J.; and Zhu, J. 2024. Accelerating transformer pre-training with 2:4 sparsity. In *ICML*.
- Hubara, I.; Chmiel, B.; Island, M.; Banner, R.; Naor, J.; and Soudry, D. 2021. Accelerated sparse neural training: A provable and efficient method to find n: m transposable masks. In *NeurIPS*, 21099–21111.
- Jiao, X.; Yin, Y.; Shang, L.; Jiang, X.; Chen, X.; Li, L.; Wang, F.; and Liu, Q. 2020. Tinybert: Distilling bert for natural language understanding. In *EMNLP*, 4163–4174.
- Kullback, S.; and Leibler, R. A. 1951. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1): 79–86.
- Kurtic, E.; and Alistarh, D. 2022. GMP*: Well-Tuned Gradual Magnitude Pruning Can Outperform Most BERT-Pruning Methods. *arXiv preprint arXiv:2210.06384*.
- Kurtic, E.; Kuznedelev, D.; Frantar, E.; Goin, M.; and Alistarh, D. 2023. Sparse finetuning for inference acceleration of large language models. In *NeurIPS*.
- LeCun, Y.; Denker, J.; and Solla, S. 1990. Optimal brain damage. In *NeurIPS*, 598–605.
- Lee, N.; Ajanthan, T.; and Torr, P. H. 2019. SNIP: Single-shot network pruning based on connection sensitivity. In *ICLR*.
- Lin, J.; Tang, J.; Tang, H.; Yang, S.; Dang, X.; and Han, S. 2024. Awq: Activation-aware weight quantization for LLM compression and acceleration. In *MLSys*, 87–100.
- Liu, Z.; Li, J.; Shen, Z.; Huang, G.; Yan, S.; and Zhang, C. 2017. Learning efficient convolutional networks through network slimming. In *ICCV*, 2736–2744.
- Mihaylov, T.; Clark, P.; Khot, T.; and Sabharwal, A. 2018. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *EMNLP*, 2381–2391.
- Molchanov, P.; Mallya, A.; Tyree, S.; Frosio, I.; and Kautz, J. 2019. Importance estimation for neural network pruning. In *CVPR*, 11264–11272.

- Nova, A.; Dai, H.; and Schuurmans, D. 2023. Gradient-free structured pruning with unlabeled data. In *ICML*, 26326–26341.
- Nowak, A. I.; Grooten, B.; Mocanu, D. C.; and Tabor, J. 2024. Compressing LLMs: The truth is rarely pure and never simple. In *ICLR*.
- Paul, M.; Chen, F.; Larsen, B. W.; Frankle, J.; Ganguli, S.; and Dziugaite, G. K. 2023. Unmasking the lottery ticket hypothesis: What’s encoded in a winning ticket’s mask? In *ICLR*.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140): 1–67.
- Renda, A.; Frankle, J.; and Carbin, M. 2020. Comparing rewinding and fine-tuning in neural network pruning. In *ICLR*.
- Sakaguchi, K.; Bras, R. L.; Bhagavatula, C.; and Choi, Y. 2021. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9): 99–106.
- Shen, M.; Yin, H.; Molchanov, P.; Mao, L.; Liu, J.; and Alvarez, J. M. 2022. Structural pruning via latency-saliency knapsack. *NeurIPS*, 12894–12908.
- Shi, D.; Tao, C.; Jin, Y.; Yang, Z.; Yuan, C.; and Wang, J. 2023. Upop: Unified and progressive pruning for compressing vision-language transformers. In *ICML*, 31292–31311.
- Singh, S. P.; and Alistarh, D. 2020. Woodfisher: Efficient second-order approximation for neural network compression. In *NeurIPS*, 18098–18109.
- Sun, M.; Liu, Z.; Bair, A.; and Kolter, J. Z. 2023. A simple and effective pruning approach for large language models. In *ICML*.
- Sun, Z.; Yu, H.; Song, X.; Liu, R.; Yang, Y.; and Zhou, D. 2020. Mobilebert: A compact task-agnostic bert for resource-limited devices. In *ACL*, 2158–2170.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *ICLR*.
- Wang, Z.; Wohlwend, J.; and Lei, T. 2019. Structured pruning of large language models. In *EMNLP*, 6151–6162.
- Wei, J.; Tay, Y.; Bommasani, R.; Raffel, C.; Zoph, B.; Borgeaud, S.; Yogatama, D.; Bosma, M.; Zhou, D.; Metzler, D.; Chi, E. H.; Hashimoto, T.; Vinyals, O.; Liang, P.; Dean, J.; and Fedus, W. 2022. Emergent abilities of large language models. In *TMLR*.
- Xia, M.; Gao, T.; Zeng, Z.; and Chen, D. 2024. Sheared llama: Accelerating language model pre-training via structured pruning. In *ICLR*.
- Zellers, R.; Holtzman, A.; Bisk, Y.; Farhadi, A.; and Choi, Y. 2019. HellaSwag: Can a machine really finish your sentence? In *ACL*, 4791–4800.
- Zhang, S.; Roller, S.; Goyal, N.; Artetxe, M.; Chen, M.; Chen, S.; Dewan, C.; Diab, M.; Li, X.; Lin, X. V.; Mi-haylov, T.; Ott, M.; Shleifer, S.; Shuster, K.; Simig, D.; Koura, P. S.; Sridhar, A.; Wang, T.; and Zettlemoyer, L. 2022. OPT: Open Pre-trained Transformer Language Models. *arXiv:2205.01068*.
- Zhang, Y.; Bai, H.; Lin, H.; Zhao, J.; Hou, L.; and Cannistraci, C. V. 2024a. Plug-and-play: An efficient post-training pruning method for large language models. In *ICLR*.
- Zhang, Y.; Zhao, L.; Lin, M.; Sun, Y.; Yao, Y.; Han, X.; Tanner, J.; Liu, S.; and Ji, R. 2024b. Dynamic sparse no training: Training-free fine-tuning for sparse llms. *ICLR*.
- Zhou, A.; Ma, Y.; Zhu, J.; Liu, J.; Zhang, Z.; Yuan, K.; Sun, W.; and Li, H. 2021. Learning n:m fine-grained structured sparse neural networks from scratch. In *ICLR*.
- Zhou, Y.; Yang, Y.; Chang, A.; and Mahoney, M. W. 2023. A three-regime model of network pruning. In *ICML*, 42790–42809.
- Zhu, M.; and Gupta, S. 2018. To prune, or not to prune: Exploring the efficacy of pruning for model compression. In *ICLR*.