

# HIGH-FIDELITY PRUNING FOR LARGE LANGUAGE MODELS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Large Language Models (LLMs) have demonstrated exceptional performance across a wide range of tasks, yet their significant computational and memory requirements present major challenges for deployment. A common approach uses Taylor expansion on the loss function to estimate neuron importance. However, its reliance on one-hot cross entropy loss, a key limitation is that it narrowly assesses importance based only on the probability assigned to the single predicted next token, thereby ignoring the other potential predictions of the original model. An intuitive solution to address this is to employ self distillation criterion for importance evaluation. However, this approach introduces significant computational overhead by requiring a separate teacher model for supervision. To this end, we propose a simple but effective criterion, information entropy of the model’s output distribution, to efficiently evaluate importance scores of neurons with Taylor pruning without requirement of additional teacher. Compared to plain cross entropy criterion, it provides a more holistic criterion for Taylor pruning to prune neurons with the least impact on the prediction of model in a global manner, thereby preserving the fidelity of the model’s predictive capabilities. Experimental results on extensive zero-shot benchmarks demonstrate that our method consistently outperforms existing pruning methods across the LLaMA and Qwen series models. The source code and trained weights will be publicly available.

## 1 INTRODUCTION

Large Language Models have become foundational in modern artificial intelligence (OpenAI et al., 2024; Touvron et al., 2023; Bai et al., 2023; DeepSeek-AI et al., 2024). This success, however, is built upon models with billions of parameters, whose substantial computational and memory footprints present significant barriers to their widespread deployment, particularly in resource-constrained environments. To reduce computational and memory footprints of LLMs, we analyze the proper components for pruning from two aspects, the proportion of parameters and the operational risk of performance degradation. We find that Multi-Layer Perceptron (MLP) modules constitute the vast majority of parameters in modern LLMs, thus offering the most significant opportunity for parameter reduction (Kaplan et al., 2020). For instance, in the Llama2-7B model, MLP modules account for approximately 68.3% of the parameters of the whole model. Moreover, we observe that pruning attention heads is a coarse-grained operation, akin to removing entire functional units (Voita et al., 2019), which poses a high risk of catastrophic performance degradation. Therefore, in this paper, we focus on the pruning of MLP modules to facilitate a more stable trade-off between model capacity and efficiency.

To minimize performance loss after pruning, Taylor-based pruning methods (Molchanov et al., 2017; 2016; 2019) estimate parameter importance through its first-order influence on a given loss function using Taylor expansion and prune the parameters with the least importance. However, as shown in Figure 1 (a), the efficacy of these methods is fundamentally hampered by their reliance on one-hot cross entropy criterion, which measures how accurately the model predicts a single, ground-truth next token. This criterion is designed to assess importance based only on the single ground-truth next token, while ignoring all other potential predictions. Consequently, the pruning process guided by this criterion only minimizes the change of label-related prediction, failing to preserve the rich knowledge encoded across the model’s full output distribution.

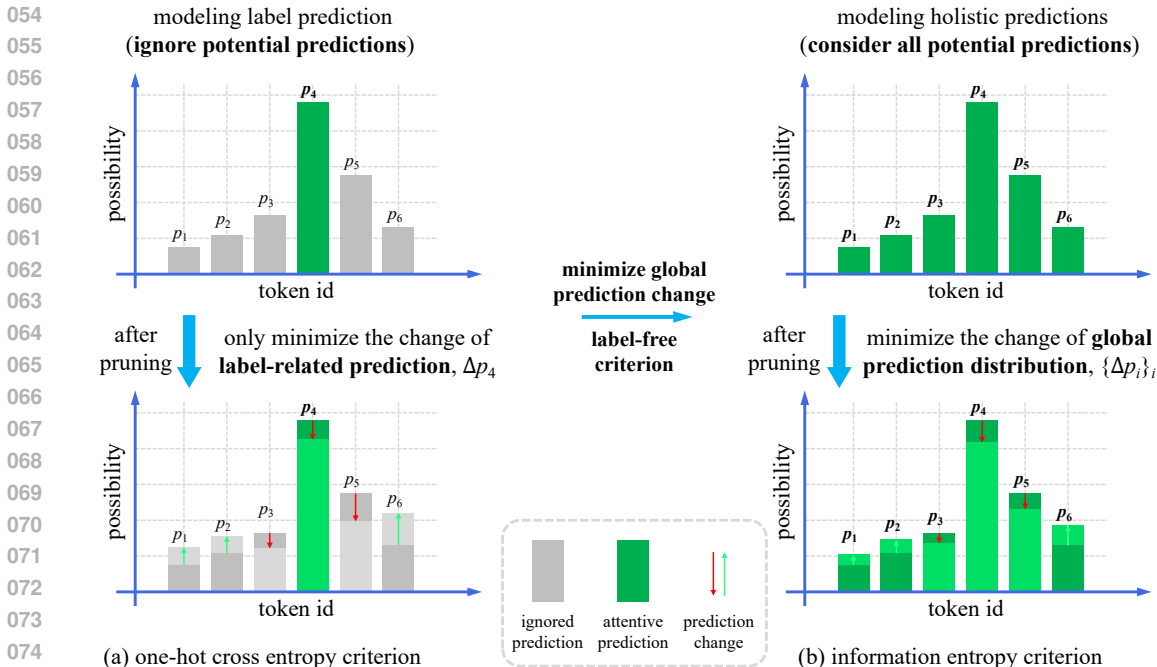


Figure 1: One-hot label loss criterion vs our proposed information entropy criterion. The cross entropy criterion (**left**) measures the model prediction by only label-related prediction. Based on this criterion, Taylor-based pruning only minimizes the change of label-related prediction after pruning. In contrast, our proposed information entropy (**right**) fully represent holistic predictions, thus minimizing the change of global prediction distribution for better performance preservation.

An intuitive solution is to employ self distillation criterion for importance evaluation as SDMP-Prune (Zhu & Shen, 2025). However, this approach suffers from two significant drawbacks. It not only introduces significant computational overhead by requiring a separate teacher model for supervision, but it also suffers from a more critical defect, where the initial distillation loss is zero, leaving no gradient to guide the initial importance scoring.

To address the above issues, we propose a novel criterion for Taylor based neuron importance evaluation. Instead of conventional loss function, we introduce information entropy of the model’s prediction distribution as the core criterion for Taylor based importance evaluation as Figure 1 (b). This shift in criterion offers two distinct advantages over existing methods. First, unlike the narrow focus of cross entropy on a single target token, our criterion can model holistic distribution of potential predictions. It provides a more faithful measure of neurons’ contribution by considering all potential predictions, ensuring the pruning process aims to minimize the change of the global prediction distribution. Second, our method is conceptually simpler and more efficient than self distillation approaches like SDMP-Prune. It circumvents the need for a separate teacher model, eliminating the associated computational overhead and avoiding the critical flaw of a null initial gradient. Our complete process involves pruning neurons with the lowest entropy-based importance scores, followed by a brief fine-tuning phase to efficiently restore the model’s performance.

Overall, our main contributions can be summarized as follows.

- We introduce a novel pruning criterion based on information entropy for Taylor-based pruning, which creates an elegant, efficient, and fundamentally label-free criterion for importance scoring.
- By modeling holistic predictions, our entropy-based criterion results in a more accurate importance assessment, which better preserves the model’s intrinsic knowledge after pruning by minimizing the change of the global prediction distribution.
- Extensive zero-shot benchmarks demonstrate that our method consistently outperforms existing techniques on the LLaMA and Qwen series models. Especially for LLaMA2-7B,

with 20% parameters and FLOPs reduction, our pruned model not only recovers but even exceed the performance of the original dense model after a brief fine-tuning.

## 2 RELATED WORK

**Foundational Importance Metrics.** Early pruning methods use static metrics to assess parameter importance. Gradient-based approaches employ Taylor expansions, from first-order methods (Molchanov et al., 2017; Ma et al., 2023) to second-order Hessian-based improvements (Molchanov et al., 2019). These gradient-based methods typically adopt one-hot cross entropy criterion, modeling label prediction while ignoring potential predictions. In contrast, our information entropy criterion models holistic predictions by considering all potential predictions. LoRAPrune (Zhang et al., 2023) reduces costs by using LoRA gradients as importance proxies. Gradient-free alternatives include Wanda (Sun et al., 2023), which combines weight magnitudes with activation norms, and FLAP (An et al., 2024), which uses activation fluctuations. Reconstruction-based methods like SlimGPT (Ling et al., 2024) minimize post-pruning squared error, while SparseGPT (Frantar & Alistarh, 2023) frames pruning as layer-wise sparse regression for efficient one-shot pruning.

**Pruning Strategies.** Some methods employ dynamic sparsity allocation strategies. OWL (Yin et al., 2023) sets layer-specific sparsity based on outlier weight distributions, while SlimLLM (Guo et al., 2025) uses input-output cosine similarity for pruning ratios. Layer removal approaches like ShortGPT (Men et al., 2024) use “Block Influence” metrics to discard redundant layers, while Chen et al. (2025) replace such layers with lightweight MLPs. Other methods include APT (Zhao et al., 2024), which allocates rates based on layer sensitivity, and learnable approaches like Compresso (Guo et al., 2023) with mask generators and LoRAP (Li et al., 2024) with differentiated FFN/MHA strategies. After pruning, these approaches only minimize the change of label-related prediction, whereas our method minimizes the change of global prediction distribution.

**Combination with Other Techniques and Performance Recovery.** Structured pruning requires performance recovery to mitigate accuracy degradation. Existing methods present distinct trade-offs: Fine-tuning directly retrains pruned models, achieving reliable accuracy restoration but with high computational cost. Low-rank integration (Hsu et al., 2022; Yu & Wu, 2023) and knowledge distillation (Muralidharan et al., 2024) offer parameter efficiency and guided learning respectively, yet typically underperform compared to fine-tuning. Retraining-free approaches like Olica (He & Lin, 2025) use linear calibration for computational efficiency but provide limited recovery.

**Entropy-Based Pruning Criterion.** Some methods use information entropy as pruning criterion. NEPENTHE (Liao et al., 2024) removes weights from low-entropy activation layers, while DenoiseRotator (Gu et al., 2025) reshapes importance distributions through learned transformations, concentrating scores on fewer weights to minimize entropy and improve pruning effectiveness. Our information entropy criterion distinguishes itself by modeling holistic predictions and considering all potential predictions, enabling us to minimize the change of global prediction distribution after pruning with the advantage of a label-free criterion.

## 3 PRELIMINARY

In modern Transformer-based LLMs, MLP module within each block has evolved, often taking the form of a SwiGLU network.

$$\text{MLP}(x) = (\text{SiLU}(xW_{\text{gate}}) \odot (xW_{\text{up}})) W_{\text{down}}, \quad (1)$$

where the layer is configured without the bias term. The input is  $x \in \mathbb{R}^{d_{\text{model}}}$ . The weight matrices  $W_{\text{gate}}, W_{\text{up}} \in \mathbb{R}^{d_{\text{model}} \times d_{\text{hidden}}}$ , and  $W_{\text{down}} \in \mathbb{R}^{d_{\text{hidden}} \times d_{\text{model}}}$ . The hidden activation  $h = \text{SiLU}(xW_{\text{gate}}) \odot (xW_{\text{up}})$  has dimensions  $h \in \mathbb{R}^{d_{\text{hidden}}}$ , and its individual components  $h_i$  represent the hidden neurons, which we aim to prune.

To determine which neurons to prune, we estimate their importance by the effect their removal would have on the loss of model predictions,  $\mathcal{L}$ . We efficiently evaluate the effect of neuron removal using

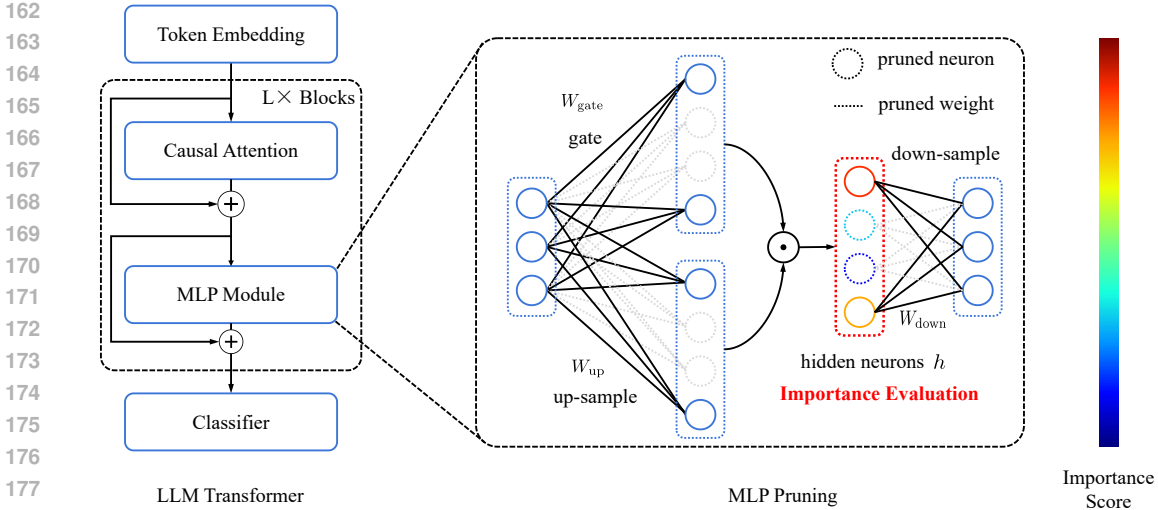


Figure 2: The overview of our high fidelity prune method. We focus on the pruning of hidden neurons  $h$  of parameter-intensive MLP modules of LLM Transformer. To this end, we apply information entropy of model prediction to evaluate the importance scores of hidden neurons  $h$  based on Taylor expansion. Then, we rank the hidden neurons  $h$  by the obtained importance scores and prune the least important neurons, thus reducing model size while minimizing performance degradation.

first-order Taylor expansion. The change in loss,  $\Delta\mathcal{L}_i$ , from ablating the  $i$ -th neuron (i.e., setting its activation  $h_i$  to zero) can be approximated as

$$\Delta\mathcal{L}_i = \mathcal{L}(h_i = 0) - \mathcal{L}(h_i) \approx \frac{\partial\mathcal{L}}{\partial h_i}(0 - h_i) = -\frac{\partial\mathcal{L}}{\partial h_i}h_i. \quad (2)$$

The magnitude of loss change,  $\mathcal{I}(h_i) = |\Delta\mathcal{L}_i| = \left| \frac{\partial\mathcal{L}}{\partial h_i}h_i \right|$ , is then used as the importance score of neuron  $h_i$ .

## 4 METHOD

### 4.1 OVERVIEW

In this paper, we propose HFPrune, a novel method that compress large language models by structurally pruning the Multi-Layer Perceptron (MLP) modules within their transformer architecture. The HFPrune process consists of three main stages. First, we establish a novel, label-free criterion to quantify neuron importance, the information entropy of the model’s global prediction distribution as shown in Figure 1 (b). Unlike previous approaches that rely on one-hot cross entropy criterion and focus exclusively on modeling label prediction while ignoring all potential predictions, our method models holistic predictions by considering all potential predictions across the entire vocabulary. Then, we compute an importance score for each neuron by applying first-order Taylor expansion to this entropy-based criterion. Second, based on these scores, we prune the model by removing a fixed ratio of the least influential neurons from each MLP layer, as depicted in Figure 2. After pruning, while previous methods only minimize the change of label-related predictions, our approach minimizes the change of the global prediction distribution. Finally, a brief fine-tuning phase is employed to counteract any performance degradation and restore the model’s capabilities.

### 4.2 INFORMATION ENTROPY FOR IMPORTANCE EVALUATION

Diverging from previous Taylor-based methods that rely on one-hot cross entropy criterion and focus solely on label prediction, our approach introduces an information entropy criterion for assessing neuron importance. Specifically, for a given input  $x$ , we define our criterion  $\mathcal{C}_H(x)$ , as the information entropy of the model’s prediction distribution  $P = \{p_1(x), p_2(x), \dots, p_V(x)\}$  over a

**Algorithm 1** High Fidelity Pruning

---

```

216 1: Input: Pre-trained model  $M$ , calibration dataset  $\mathcal{D}_{\text{calib}}$ , MLP pruning ratio  $\rho_{\text{mlp}}$ .
217 2: Output: Pruned model  $M'$ .
218 3: for each MLP module  $l$  in  $M$  do
219 4:   Initialize importance scores  $\{\mathcal{I}_i\}_i = 0$  for each neuron  $i$  in module  $l$ .
220 5:   for each input  $x \in \mathcal{D}_{\text{calib}}$  do
221 6:     Compute hidden activations  $h(x)$  of MLP module  $l$ .
222 7:     Compute the information entropy of the model prediction  $\mathcal{C}_H(x)$ .
223 8:     Compute gradients  $\nabla_{h(x)} \mathcal{C}_H(x) = \frac{\partial \mathcal{C}_H(x)}{\partial h(x)}$  by backward computation.
224 9:     for each neuron  $i$  in hidden layer of MLP module  $l$  do
225 10:       $\mathcal{I}_i \leftarrow \mathcal{I}_i + \left| \frac{\partial \mathcal{C}_H(x)}{\partial h_i(x)} h_i(x) \right|$ .
226 11:     end for
227 12:   end for
228 13:   for each neuron  $i$  in MLP module  $l$  do
229 14:      $\mathcal{I}_i \leftarrow \frac{\mathcal{I}_i}{|\mathcal{D}_{\text{calib}}|}$ .
230 15:   end for
231 16:   Let  $k = \lfloor \rho_{\text{mlp}} \cdot d_{\text{hidden}} \rfloor$ .
232 17:   Identify the set  $K$  of indices of the  $k$  neurons with the lowest scores in  $\mathcal{I}$ .
233 18:   Update MLP module  $l$  by removing neurons with indices in  $K$ .
234 19:   Obtain the final model  $M'$ .
235 20: end for

```

---

vocabulary size of  $V$  as

$$\mathcal{C}_H(x) = - \sum_{j=1}^V p_j(x) \log_2 p_j(x). \quad (3)$$

This information entropy criterion provides a holistic signal capturing the model’s predictive confidence across the entire vocabulary, considering all potential predictions rather than focusing on specific labels. Moreover, our proposed information entropy criterion doesn’t rely on the given labels for training, thus allowing label-free Taylor pruning. We now apply it on Taylor-based importance evaluation to quantify the importance of each hidden neuron. For input  $x$ , the importance of the  $i$ -th neuron,  $\mathcal{I}_i(x)$ , is the magnitude of the entropy change caused by ablating its activation,  $h_i(x)$ . To obtain a robust and generalized score, we average this value across a calibration dataset,  $\mathcal{D}_{\text{calib}}$ , to compute the final importance score,  $\mathcal{I}_i(x)$  as

$$\mathcal{I}_i(x) = \frac{1}{|\mathcal{D}_{\text{calib}}|} \sum_{x \in \mathcal{D}_{\text{calib}}} \left| \frac{\partial \mathcal{C}_H(x)}{\partial h_i(x)} h_i(x) \right|. \quad (4)$$

A higher score  $\mathcal{I}_i$  signifies that the neuron is more critical for preserving the fidelity of the model’s global prediction distribution.

### 4.3 HIGH FIDELITY PRUNING

With the importance score  $\mathcal{I}_i$  computed for every hidden neuron, we proceed to the structural pruning phase. This process is applied uniformly across all MLP layers. For each layer, we rank its hidden neurons by their importance scores and remove a fraction of those deemed least influential the ones with the lowest scores. The fraction of neurons to be pruned is controlled by a predefined ratio,  $\rho_{\text{mlp}}$ . Specifically, for a layer with  $d_{\text{hidden}}$  neurons, we remove the  $k = \lfloor \rho_{\text{mlp}} \cdot d_{\text{hidden}} \rfloor$  neurons with the lowest importance scores. This removal is achieved by modifying the weight matrices. Specifically, for each pruned neuron, we excise the corresponding row from the up-projection  $W_{\text{up}}$  and gate-projection  $W_{\text{gate}}$  matrices and the corresponding column from the down-projection  $W_{\text{down}}$ .

By pruning based on our information entropy criterion, we ensure that the resulting compressed model minimizes the change in global prediction distribution rather than just preserving label-specific predictions. This comprehensive approach maintains the model’s distributional integrity across all potential outputs, leading to better preservation of the model’s overall predictive capabilities. The overall algorithm can be summarized as Algorithm 1.

Table 1: Zero-shot performance on LLaMA-2-7B, which are finetuned on the LaMini dataset. The average accuracy is calculated among 10 zero-shot benchmarks. The “**bolded**” marked results denotes the best one under the same pruning ratio.

Ratio	Method	ARCc	ARCe	BoolQ	Crows	OBQA	PIQA	Race	SiQA	TfQA	Wino	Average
0%	Llama-2-7B	45.1	73.8	79.4	67.4	44.2	78.7	40.1	46.5	38.8	69.3	58.3
20%	LLM-pruner	40.4	70.1	80.2	61.7	38.8	75.8	39.0	47.1	43.9	64.3	56.1
	LoRAPrune	41.6	71.0	81.7	58.7	41.4	76.7	40.4	44.0	<b>65.9</b>	65.9	56.7
	LoRAP	38.5	66.0	70.9	–	39.6	<b>78.1</b>	–	–	–	65.7	–
	SDMPPrune	43.9	72.3	81.7	<b>62.1</b>	42.0	77.0	41.3	48.5	44.9	<b>68.4</b>	58.2
	HFPrune (Ours)	<b>47.1</b>	<b>73.8</b>	<b>85.2</b>	60.2	<b>43.2</b>	77.3	<b>43.3</b>	<b>49.5</b>	44.7	66.2	<b>59.0</b>
30%	LLM-pruner	38.0	64.8	75.6	<b>62.3</b>	36.4	73.4	35.7	47.3	42.3	62.9	53.9
	LoRAPrune	38.6	65.1	74.1	61.4	37.4	72.9	39.0	46.3	<b>44.8</b>	<b>66.5</b>	54.6
	LoRAP	35.5	60.6	69.6	–	37.8	<b>76.7</b>	–	–	–	63.0	–
	SDMPPrune	39.6	67.9	80.4	58.5	37.2	75.2	<b>40.0</b>	47.8	43.7	65.4	55.6
	HFPrune (Ours)	<b>41.9</b>	<b>70.2</b>	<b>82.9</b>	58.1	<b>40.0</b>	75.2	39.5	<b>48.8</b>	44.2	62.4	<b>56.3</b>

## 5 EXPERIMENT

### 5.1 EXPERIMENTAL SETUP

**Evaluation and Datasets.** To evaluate the effectiveness our method, we evaluate our proposed HFPrune on LLaMA and Qwen series models. Following prior work (Ma et al., 2023; Zhang et al., 2024), we assess the zero-shot performance of the pruned models using the lm\_eval (Gao et al., 2021). Our evaluation encompasses a diverse set of zero-shot benchmarks, including ARC-easy (Clark et al., 2018), ARC-challenge (Clark et al., 2018), BOOLQ (Clark et al., 2019), Crows-Pairs (Nangia et al., 2020), OpenBookQA (Mihaylov et al., 2018), PIQA (Bisk et al., 2020), Race (Lai et al., 2017), SocialQA (Sap et al., 2019), TruthfulQA (Lin et al., 2021), and Winogrande (Sakaguchi et al., 2020).

**Calibration and Fine-Tuning Datasets.** We utilize two distinct datasets for the calibration and fine-tuning stage, respectively. In the calibration stage, we obtain importance score using a calibration dataset of 43,128 sequences randomly sampled from the C4 dataset (Raffel et al., 2020). Following the protocol of Wanda (Sun et al., 2023), each sequence is processed to a fixed length of 1,024 tokens, where longer sequences are randomly cropped into 1,024 tokens. In the fine-tuning stage, we use the LaMini-instruction dataset (Wu et al., 2024) across all experiments for fair comparison. This dataset consists of instruction-response pairs generated by GPT-3.5-turbo from various prompt resources.

**Implementation Details.** To recover performance after pruning, each model variant undergoes a brief fine-tuning stage. We fine-tune each model for 2 epochs on the LaMini-instruction dataset using LoRA strategy. This process employs the AdamW optimizer with BF16 precision and a cosine learning rate scheduler. Further implementation details, including specific hyperparameters, are available in the Section A.1 of appendix.

### 5.2 EXPERIMENTAL RESULTS

#### 5.2.1 COMPARISON OF EXISTING METHODS

To facilitate a fair comparison against existing methods, we apply HFPrune to models from the LLaMA and Qwen series. This process yields a suite of model variants at different sparsity levels, enabling a thorough evaluation of the trade-off between compression and performance.

To validate the effectiveness of our method, we compare our method with existing structural pruning methods on LLaMA-2-7B model and then evaluate the pruned model on 10 popular zero-shot benchmarks. The results are listed in Table 1. The results demonstrate that our method significantly outperforms other structural pruning methods, when 20% and 30% parameters of LLaMA-2-7B model are pruned. When pruning ratio is set to 20%, our method achieves 59% average accuracy over 10 popular zero-shot benchmarks, surpassing the second best method, SDMPPrune, by 0.8%. Notably, our method even outperforms the original model by 0.7%. We further explore high pruned

Table 2: Zero-shot performance on smaller LLMs, including LLaMA3.2-3.2B (left) and LLaMA3.2-1.2B (right) models. The settings are identical to the ones of LLaMA-2-7B. More detailed results on each benchmark are available in the Appendix.

(a) LLaMA3.2-3.2B Results				(b) LLaMA3.2-1.2B Results			
Prune%	Method	#Params	Avg Zero-shot A. $\uparrow$	Prune%	Method	#Params	Avg Zero-shot A. $\uparrow$
-	-	3.2B	54.30	-	-	1.2B	51.47
20%	LLMPruner	2.56B	52.24	20%	LLMPruner	0.989B	47.39
	LoRAPrune	2.57B	50.30		LoRAPrune	0.989B	47.17
	SDMPPrune	2.56B	53.37		SDMPPrune	0.989B	48.94
	<b>HFPrune(Ours)</b>	<b>2.56B</b>	<b>54.07</b>		<b>HFPrune(Ours)</b>	<b>0.989B</b>	<b>50.77</b>
30%	LLMPruner	2.25B	48.68	30%	LLMPruner	0.865B	44.79
	LoRAPrune	2.25B	48.36		LoRAPrune	0.865B	44.86
	SDMPPrune	2.25B	51.03		SDMPPrune	0.865B	46.40
	<b>HFPrune(Ours)</b>	<b>2.25B</b>	<b>52.28</b>		<b>HFPrune(Ours)</b>	<b>0.865B</b>	<b>48.59</b>

ing ratio as 30%, our HFPrune is consistently superior to other methods. The above results indeed support the superiority of our method on LLM compression task.

To further verify the generalization on smaller scale LLMs, our proposed HFPrune is also applied on LLaMA3.2-3.2B and LLaMA3.2-1.2B models. The results shown in Table 2 also demonstrate that our method can effectively reduce the performance gap between the pruned model and the original model.

Table 3: Zero-shot performance on Qwen series models. The settings are identical to the ones of LLaMA-2-7B.

Model	Ratio	Method	ARCc	ARCe	BoolQ	Crows	OBQA	PIQA	Race	SiQA	TfQA	Wino	Average
Qwen2.5-7B	0%	Original	51.1	77.5	84.7	64.0	46.8	79.8	41.4	54.8	56.4	73.2	63.0
	20%	SDMPPrune	48.1	69.7	<b>84.6</b>	<b>64.2</b>	44.4	77.6	<b>41.7</b>	48.0	51.8	<b>67.2</b>	59.7
		<b>HFPrune(Ours)</b>	<b>52.1</b>	<b>74.5</b>	<b>84.6</b>	62.9	<b>45.4</b>	<b>78.1</b>	41.2	<b>48.1</b>	<b>52.0</b>	64.8	<b>60.4</b>
	30%	SDMPPrune	40.5	63.6	81.1	<b>62.2</b>	38.4	74.2	39.4	44.5	<b>49.1</b>	60.2	55.3
		<b>HFPrune(Ours)</b>	<b>48.2</b>	<b>70.3</b>	<b>83.2</b>	60.4	<b>41.6</b>	<b>76.1</b>	<b>40.7</b>	<b>48.9</b>	48.2	<b>62.8</b>	<b>58.0</b>
	Qwen2.5-1.5B	0%	Original	45.4	72.0	73.0	60.9	40.4	76.0	36.5	48.9	46.6	63.5
20%		SDMPPrune	32.3	59.2	72.1	<b>56.2</b>	35.2	72.0	37.7	43.6	<b>44.7</b>	58.2	51.1
		<b>HFPrune(Ours)</b>	<b>41.8</b>	<b>68.8</b>	<b>79.4</b>	55.3	<b>39.4</b>	<b>74.1</b>	<b>38.7</b>	<b>46.4</b>	42.2	<b>59.8</b>	<b>54.6</b>
30%		SDMPPrune	26.3	49.5	62.9	<b>55.0</b>	32.8	68.6	34.9	42.1	43.8	57.9	47.4
		<b>HFPrune(Ours)</b>	<b>35.2</b>	<b>62.6</b>	<b>76.0</b>	54.0	<b>34.6</b>	<b>70.8</b>	<b>36.2</b>	<b>44.0</b>	<b>44.5</b>	<b>58.6</b>	<b>51.7</b>
Qwen3-1.7B		0%	Original	42.8	70.0	77.6	55.2	37.8	71.9	36.8	44.9	46.0	61.6
	20%	SDMPPrune	31.3	58.5	70.8	53.7	33.4	71.4	37.1	43.8	44.7	<b>58.6</b>	50.3
		<b>HFPrune(Ours)</b>	<b>39.1</b>	<b>69.4</b>	<b>78.9</b>	<b>55.8</b>	<b>36.2</b>	<b>72.4</b>	<b>39.7</b>	<b>46.4</b>	<b>46.4</b>	58.2	<b>54.3</b>
	30%	SDMPPrune	25.9	45.1	62.3	55.1	29.8	66.2	35.3	40.5	43.4	<b>56.7</b>	46.0
		<b>HFPrune(Ours)</b>	<b>34.5</b>	<b>61.9</b>	<b>78.2</b>	<b>55.5</b>	<b>33.8</b>	<b>70.8</b>	<b>36.3</b>	<b>44.1</b>	<b>45.2</b>	56.4	<b>51.7</b>

We also assess the effectiveness of our method on Qwen series models, including Qwen2.5-1.5B, Qwen2.5-7B and Qwen3-1.7B. For brief, we focus on the comparative experiments with the previous best methods, SDMPPrune. The results are reported in Table 3. Our method consistently surpasses SDMPPrune across various model sizes and pruning ratios. We analyze the potential reason as follows. Despite SDMPPrune aims to exploit full prediction distribution, it still relies on Taylor pruning based on one-hot cross entropy criterion in the initial stage due to zero-gradient issue of self distillation. SDMPPrune fails to capture holistic prediction distribution during importance evaluation, thus leading to inferior performance. On the contrary, our method considers all potential predictions during Taylor-based importance evaluation, thus providing more accurate pruning criterion.

### 5.2.2 ACCELERATION PERFORMANCE OF PRUNED MODELS

We evaluate the practical acceleration benefits of HFPrune on the pruned LLaMA2-7B models, with results presented in Table 4. Prefill latency was measured using an input sequence of 256 tokens and a single-token generation task. To measure decoding throughput, we set the batch size to 8 and recorded the latency for generating 256 tokens. The results demonstrate a direct correlation between pruning and efficiency. As the pruning ratio increases, both model parameters and prefill latency decrease significantly. Notably, pruning 30% of the MLP layers results in a  $1.47\times$  speedup in prefill latency, confirming a tangible inference advantage.

Table 4: Acceleration results of the structural pruned LLMs. Prefill Latency and decoding throughput were measured on a single NVIDIA A6000 GPU without employing any acceleration frameworks.

Model	#Params	Ratio	Avg zero-shot $\uparrow$	Prefill (ms) $\downarrow$	Speedup	Decoding (tokens/s) $\uparrow$
LLaMA2-7B	6.7B	0	58.3	57.5	1.00 $\times$	473.9
	5.4B	20%	59.0	46.3	1.24 $\times$	553.9 (+17.9%)
	4.7B	30%	56.3	42.1	1.35 $\times$	593.8 (+25.3%)

Beyond the inference speed of pruned model, we also assess the computational efficiency of the pruning process itself, comparing HFPrune against SDMPPrune (Zhu & Shen, 2025). As detailed in Table 5, HFPrune is substantially more efficient in both time and memory usage. For instance, when pruning LLaMA2-7B, HFPrune is approximately  $3\times$  faster than SDMPPrune and consumes 31% less peak GPU memory.

Table 5: Efficiency comparison of the pruning process. The metrics are measured using 1000 randomly sampled sequences from the C4 dataset, each with a length of 1024 tokens. All models were run on four A6000 GPUs.

Model	Method	Elapsed Time (s)	Peak Memory (GB)
LLaMA3.2-1.2B	HFPrune (Ours)	150.7	7.5
	SDMPPrune	317.1	12.9
LLaMA3.2-3.2B	HFPrune (Ours)	305.8	17.3
	SDMPPrune	825.2	25.7
LLaMA2-7B	HFPrune (Ours)	508.9	35.3
	SDMPPrune	1539.8	51.2

### 5.3 ABLATION STUDY

#### 5.3.1 THE EFFECT OF INFORMATION ENTROPY CRITERION

To isolate and validate the standalone effectiveness of our proposed neuron importance criterion, we conduct a direct comparison on the LLaMA-2-7B model. Our information entropy (IE) criterion is benchmarked against two key alternatives: the cross entropy (CE) loss criterion and the self-distillation (SD) loss criterion. Crucially, all comparisons are performed without any post-pruning fine-tuning. The results in Table 6 demonstrate clear superiority of our IE criterion. As pruning ratio increases from 20% to 30%, our method maintains relatively stable performance degradation compared to baseline methods, indicating superior robustness of our importance estimates. Our method achieves the highest average performance at both pruning ratios, outperforming the best baselines 0.5 percentage points respectively. This outperformance in a retraining-free scenario confirms our central hypothesis: by modeling holistic predictions rather than focusing solely on label-related predictions, our IE criterion provides fundamentally more accurate measures of neuron importance, yielding superior models immediately after pruning.

Table 6: Zero-Shot Performance Comparison of Different Pruning Criteria on LLaMA-2-7B (No Fine-Tuning). Where IE (Information Entropy) is our proposed criterion, while SD (Self-Distillation) and CE (Cross-Entropy) represent the baseline loss criteria.

Ratio	Criterion	ARCc	ARCe	BoolQ	Crows	OBQA	PIQA	Race	SiQA	TfQA	Wino	Avg
20%	CE Loss	<b>37.2</b>	62.3	71.3	58.9	39.4	74.5	36.9	43.1	36.3	<b>65.9</b>	52.6
	SD Loss	34.2	<b>63.6</b>	67.1	59.6	38.4	<b>75.0</b>	37.4	<b>44.0</b>	36.6	62.8	51.9
	IE (ours)	37.0	63.4	<b>71.8</b>	<b>60.0</b>	<b>40.0</b>	74.9	<b>37.7</b>	43.7	<b>37.0</b>	65.0	<b>53.1</b>
30%	CE Loss	<b>31.4</b>	49.0	63.1	54.1	32.4	69.4	32.8	40.9	37.2	<b>57.4</b>	46.8
	SD Loss	28.2	<b>50.3</b>	49.0	53.9	32.6	69.3	<b>33.6</b>	41.5	37.3	56.4	45.2
	IE (ours)	30.1	49.4	<b>65.8</b>	<b>54.6</b>	<b>33.2</b>	<b>70.2</b>	33.2	<b>41.7</b>	<b>37.5</b>	57.1	<b>47.3</b>

#### 5.3.2 EFFECT OF PRUNING CRITERION ON OUTPUT DISTRIBUTION

To quantitatively validate that our information entropy (IE) criterion better preserves the global prediction distribution than the cross-entropy (CE) criterion, we analyzed the pruned Llama2-7B model over 5,000 prompts from the C4 dataset. The results, presented in Table 7, confirm our method’s superior fidelity across two key metrics. First, our IE criterion achieves a consistently lower Jensen-Shannon (JS) Distance. While the improvement at 20% sparsity is modest, the advantage becomes

more pronounced at the aggressive 30% ratio. This trend highlights that our method’s ability to preserve the distribution’s overall shape is most critical in high-compression scenarios. Second, the Top-15 Jaccard Similarity results provide complementary evidence, showing our approach better maintains the original model’s most probable next tokens at both sparsity levels. Qualitative examples that visually illustrate these findings are provided in Appendix A.3.

Table 7: Comparison of pruning methods based on output distribution similarity.

Ratio	Criterion	JS Distance ( $\downarrow$ )	Top-15 Jaccard ( $\uparrow$ )
20%	CE Loss	0.243	0.439
	Ours (IE)	0.241	0.445
30%	CE Loss	0.362	0.588
	Ours (IE)	0.353	0.595

### 5.3.3 PRUNE DIFFERENT PARTS

To validate our hypothesis of exclusively targeting MLP modules, we compare pruning only MLP modules versus pruning both attention and MLP modules. The results in Table 8 provide decisive evidence favoring the MLP-only strategy. MLP-only pruning consistently outperforms attention&MLP pruning across both scenarios. After fine-tuning, MLP-only achieves 61.9% average performance at 20% pruning compared to 60.3% for attention&MLP pruning. This gap widens at 30% pruning, demonstrating superior scalability. The fine-tuning gains reveal different recovery patterns. MLP-only pruning shows larger improvement potential, while attention&MLP pruning exhibits more limited recovery. This suggests MLP modules contain more recoverable, distributed knowledge. Therefore, this experiment provides evidence that focusing pruning efforts on MLP modules is a more effective and robust strategy for compressing LLMs.

Table 8: Prune different parts on LLaMA-2-7B

Ratio	Method	ARCC	ARCE	BoolQ	Crows	OBQA	PIQA	Race	SiQA	Wino	Average
20%	attn&mlp w/o tune	40.0	69.0	65.1	60.0	40.0	76.0	32.5	41.0	55.0	53.2
	attn&mlp w/ tune	48.3	73.9	84.6	<b>60.5</b>	42.8	76.7	41.9	49.1	65.0	60.3
	mlp w/o tune	37.0	63.4	71.8	60.0	40.0	74.9	37.7	43.7	65.0	54.8
	mlp w/ tune	<b>50.3</b>	<b>75.2</b>	<b>85.8</b>	60.0	<b>45.2</b>	<b>77.3</b>	<b>44.5</b>	<b>50.5</b>	<b>67.9</b>	<b>61.9</b>
30%	attn&mlp w/o tune	33.0	57.0	56.0	54.0	33.4	70.0	24.2	38.1	52.3	46.4
	attn&mlp w/ tune	44.6	73.0	82.9	56.3	41.4	76.0	39.9	48.1	60.1	58.0
	mlp w/o tune	30.1	49.4	65.8	54.6	33.2	70.2	33.6	41.7	57.1	48.4
	mlp w/ tune	<b>47.6</b>	<b>73.4</b>	<b>83.9</b>	<b>60.3</b>	<b>43.2</b>	<b>76.7</b>	<b>42.2</b>	<b>49.6</b>	<b>63.2</b>	<b>60.0</b>

## 6 CONCLUSION

In this paper, we present HFPrune, a novel structured pruning method for large language models that overcomes a fundamental limitation of existing Taylor-based approaches. Our core innovation is replacing the one-hot cross entropy criterion with the information entropy of the model’s global prediction distribution as the criterion for assessing neuron importance. This entropy-based criterion offers distinct advantages. By modeling holistic predictions, it provides a more faithful measure of a neuron’s contribution to the model’s overall predictive fidelity. The pruning process is thus guided to minimize the change of the global prediction distribution, preserving intrinsic knowledge far more effectively than methods that only focus on a single target token. Furthermore, our approach avoids the computational overhead and gradient initialization issues inherent in self distillation methods while achieving superior pruning results. Looking forward, our work opens several promising avenues for research. The entropy-based importance metric could be extended to other model compression techniques, such as quantization. Additionally, future investigations could explore its application to diverse neural network architectures or the development of adaptive pruning ratios based on layer specific entropy, potentially yielding even greater gains in compression efficiency.

## REFERENCES

- 486  
487  
488 Yuchao An, Xi Zhao, Tian Yu, Ming Tang, and Jinqiao Wang. Fluctuation-based adaptive structured pruning for large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 10865–10873, 2024.
- 489  
490  
491 Jinze Bai, Shuai Bai, and et al. Qwen technical report, 2023.
- 492  
493 Yonatan Bisk, Rowan Zellers, Ronan Lebras, Jianfeng Gao, and Yejin Choi. PIQA: Reasoning about physical commonsense in natural language. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*, pp. 7432–7439, 2020.
- 494  
495  
496 Xiaodong Chen, Yuxuan Hu, Jing Zhang, Yanling Wang, Cuiping Li, and Hong Chen. Streamlining redundant layers to compress large language models, 2025. URL <https://arxiv.org/abs/2403.19135>.
- 497  
498  
499 Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. BoolQ: Exploring the surprising difficulty of natural yes/no questions. 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)*, pp. 2924–2936, 2019.
- 500  
501  
502 Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- 503  
504  
505 DeepSeek-AI, Aixin Liu, and et al. Deepseek-v3 technical report, 2024.
- 506  
507  
508 Elias Frantar and Dan Alistarh. Sparsegpt: Massive language models can be accurately pruned in one-shot, 2023. URL <https://arxiv.org/abs/2301.00774>.
- 509  
510  
511 Leo Gao, Jonathan Tow, Bidisha Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Gold, Jeffrey Hsu, Kyle Lebrun, Teven Le Scao, Livia Ran, Conglong Su, Phil Wang, and Samuel Weaver. A framework for few-shot language model evaluation. <https://doi.org/10.5281/zenodo.5371628>, 2021.
- 512  
513  
514 Tianteng Gu, Bei Liu, Bo Xiao, Ke Zeng, Jiacheng Liu, and Yanmin Qian. Denoiserotator: Enhance pruning robustness for llms via importance concentration, 2025. URL <https://arxiv.org/abs/2505.23049>.
- 515  
516  
517 Jialong Guo, Xinghao Chen, Yehui Tang, and Yunhe Wang. Slimllm: Accurate structured pruning for large language models, 2025. URL <https://arxiv.org/abs/2505.22689>.
- 518  
519  
520 Song Guo, Jiahang Xu, Li Lyna Zhang, and Mao Yang. Compresso: Structured pruning with collaborative prompting learns compact large language models, 2023. URL <https://arxiv.org/abs/2310.05015>.
- 521  
522  
523 Jiujun He and Huazhen Lin. Olica: Efficient structured pruning of large language models without retraining, 2025. URL <https://arxiv.org/abs/2506.08436>.
- 524  
525  
526 Yen-Chun Hsu, Ting Hua, Sung-Hsien Chang, Qingqing Lou, Yun-Ti Shen, and Hong-Jin Jin. Language model compression with weighted low-rank factorization. *arXiv preprint arXiv:2207.00112*, 2022.
- 527  
528  
529 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- 530  
531  
532 Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. RACE: Large-scale ReAding comprehension dataset from examinations. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 785–794, 2017.
- 533  
534  
535 Guanru Li, Yue Tang, and Wei Zhang. Lorap: Transformer sub-layers deserve differentiated structured compression for large language models. *arXiv preprint arXiv:2404.09695*, 2024.
- 536  
537  
538  
539

- 540 Zhu Liao, Victor Quétu, Van-Tam Nguyen, and Enzo Tartaglione. Nepenthe: Entropy-based pruning  
541 as a neural network depth’s reducer, 2024. URL <https://arxiv.org/abs/2404.16890>.
- 542
- 543 Stephanie Lin, Jacob Hilton, and Owain Evans. TruthfulQA: Measuring how models mimic human  
544 falsehoods. *arXiv preprint arXiv:2109.07958*, 2021.
- 545
- 546 Guangji Ling, Zhaofan Wang, Yuzhang Yan, and Qing Liu. Slimgpt: Layer-wise structured pruning  
547 for large language models. *arXiv preprint arXiv:2412.18110*, 2024.
- 548
- 549 Xinyu Ma, Guangxin Fang, and Xin Wang. Llm-pruner: On the structural pruning of large language  
550 models. *Advances in neural information processing systems*, 36:21702–21720, 2023.
- 551
- 552 Xin Men, Mingyu Xu, Qingyu Zhang, Bingning Wang, Hongyu Lin, Yaojie Lu, Xianpei Han, and  
553 Weipeng Chen. Shortgpt: Layers in large language models are more redundant than you expect,  
2024. URL <https://arxiv.org/abs/2403.03853>.
- 554
- 555 Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct elec-  
556 tricity? a new dataset for open book question answering. In *Proceedings of the 2018 Conference*  
557 *on Empirical Methods in Natural Language Processing*, pp. 2381–2391, 2018.
- 558
- 559 Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional  
560 neural networks for resource efficient inference. *arXiv preprint arXiv:1611.06440*, 2016.
- 561
- 562 Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional  
563 neural networks for resource efficient inference, 2017. URL [https://arxiv.org/abs/](https://arxiv.org/abs/1611.06440)  
1611.06440.
- 564
- 565 Pavlo Molchanov, Arun Mallya, Stephen Tyree, Iuri Frosio, and Jan Kautz. Importance estimation  
566 for neural network pruning. In *Proceedings of the IEEE/CVF conference on computer vision and*  
*pattern recognition*, pp. 11264–11272, 2019.
- 567
- 568 Sreepathroopakumari Muralidharan, S M Turuvekere Sreenivas, Rahul Joshi, Maciej Chochowski,  
569 Md A. Patwary, Mohammad Shoeybi, Bryan Catanzaro, Jan Kautz, and Pavlo Molchanov. Com-  
570 pact language models via pruning and knowledge distillation. *Advances in Neural Information*  
*Processing Systems*, 37:41076–41102, 2024.
- 571
- 572 Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. Crows-Pairs: A challenge  
573 dataset for measuring social biases in masked language models. In *Proceedings of the 2020*  
574 *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1953–1967,  
575 2020.
- 576
- 577 OpenAI, Josh Achiam, Steven Adler, and et al. Gpt-4 technical report, 2024.
- 578
- 579 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi  
580 Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-  
581 text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. URL [http://](http://jmlr.org/papers/v21/20-074.html)  
[jmlr.org/papers/v21/20-074.html](http://jmlr.org/papers/v21/20-074.html).
- 582
- 583 Keisuke Sakaguchi, Ronan Lebras, Benjamin van Durme, and Yejin Choi. WinoGrande: An adver-  
584 sarial winograd schema challenge at scale. In *Proceedings of the AAAI Conference on Artificial*  
*Intelligence*, volume 34, pp. 8732–8740, 2020.
- 585
- 586 Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Lebras, and Yejin Choi. SocialIQA: Common-  
587 sense reasoning about social interactions. In *Proceedings of the 2019 Conference on Empirical*  
588 *Methods in Natural Language Processing and the 9th International Joint Conference on Natural*  
589 *Language Processing (EMNLP-IJCNLP)*, pp. 4906–4916, 2019.
- 590
- 591 Mingjie Sun, Zhuang Liu, Anna Bair, and J. Zico Kolter. A simple and effective pruning approach  
592 for large language models. *arXiv preprint arXiv:2306.11695*, 2023.
- 593
- 594 Hugo Touvron, Thibaut Lavril, and et al. Llama: Open and efficient foundation language models,  
2023.

- 594 Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. Analyzing multi-head  
595 self-attention: Specialized heads do the heavy lifting, the rest can be pruned, 2019. URL <https://arxiv.org/abs/1905.09418>.  
596  
597
- 598 Minghao Wu, Abdul Waheed, Chiyu Zhang, Muhammad Abdul-Mageed, and Alham Fikri Aji.  
599 LaMini-LM: A diverse herd of distilled models from large-scale instructions. In Yvette Gra-  
600 ham and Matthew Purver (eds.), *Proceedings of the 18th Conference of the European Chapter*  
601 *of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 944–964, St. Ju-  
602 lian’s, Malta, March 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.  
603 eacl-long.57. URL <https://aclanthology.org/2024.eacl-long.57/>.
- 604 Lichuan Yin, Yichi Wu, Zixuan Zhang, Chien-Yu Hsieh, Yingdong Wang, Yulong Jia, Geng Li,  
605 A. Jaiswal, M. Pechenizkiy, Yu-Chin Liang, et al. Outlier weighed layerwise sparsity (owl): A  
606 missing secret sauce for pruning llms to high sparsity. *arXiv preprint arXiv:2310.05175*, 2023.  
607
- 608 Hao Yu and Jianxin Wu. Compressing transformers: features are low-rank, but weights are not!  
609 In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 11007–11015,  
610 2023.
- 611 Mingqi Zhang, Hanchi Chen, Chuyun Shen, Zhaozhuo Yang, Liangyu Ou, Xiang Yu, and Bingbing  
612 Zhuang. Loraprune: Pruning meets low-rank parameter-efficient fine-tuning. *arXiv preprint*  
613 *arXiv:2305.18403*, 2023.
- 614 Zihan Zhang, Xinyu Ma, Jiachen Chen, Lifan Wang, Yichen Huang, Zihan Zhou, Zeyu Wang,  
615 and Yichen Liu. Loraprune: Structured pruning meets low-rank adaptation. *arXiv preprint*  
616 *arXiv:2404.09994*, 2024.
- 617 Bowen Zhao, Hannaneh Hajishirzi, and Qingqing Cao. Apt: Adaptive pruning and tuning pretrained  
618 language models for efficient training and inference, 2024. URL <https://arxiv.org/abs/2401.12200>.  
619  
620  
621
- 622 Hourun Zhu and Chengchao Shen. Sdmprune: Self-distillation mlp pruning for efficient large lan-  
623 guage models, 2025. URL <https://arxiv.org/abs/2506.11120>.  
624  
625

## 626 A APPENDIX

### 627 A.1 IMPLEMENTATION DETAILS.

628 To restore model performance post-pruning, we conduct a brief fine-tuning phase using the AdamW  
629 optimizer and BF16 precision. As detailed in Table 9, we customized the hyperparameters for  
630 each model and sparsity ratio. For instance, larger models like LLaMA2-7B and Qwen2.5-7B were  
631 trained with a larger batch size of 512 on 8×A6000 GPUs, while smaller models used a batch size  
632 of 256 on 4×4090 GPUs. The learning rates were also adjusted, ranging from  $2 \times 10^{-4}$  to  $4 \times 10^{-4}$   
633 depending on the model scale.  
634  
635

636 Table 9: Implementation Details of the structural pruned LLMs.  
637

Model	Ratio	$\rho_{mlp}$	Pack	lr	beta1	beta2	Weight decay	Precision	Batch size	Device
LLaMA3.2-1.2B	20%	0.70	yes	$4 \times 10^{-4}$	0.9	0.999	0	BF16	256	4×4090
	30%	0.55	yes	$4 \times 10^{-4}$	0.9	0.999	0	BF16	256	4×4090
LLaMA3.2-3.2B	20%	0.70	yes	$4 \times 10^{-4}$	0.9	0.999	0	BF16	256	4×4090
	30%	0.55	yes	$4 \times 10^{-4}$	0.9	0.999	0	BF16	256	4×4090
LLaMA2-7B	20%	0.70	yes	$2 \times 10^{-4}$	0.9	0.999	0	BF16	512	8×A6000
	30%	0.54	yes	$2 \times 10^{-4}$	0.9	0.999	0	BF16	512	8×A6000
Qwen2.5-1.5B	20%	0.74	yes	$4 \times 10^{-4}$	0.9	0.999	0	BF16	256	4×4090
	30%	0.60	yes	$4 \times 10^{-4}$	0.9	0.999	0	BF16	256	4×4090
Qwen3-1.7B	20%	0.68	yes	$4 \times 10^{-4}$	0.9	0.999	0	BF16	256	4×4090
	30%	0.52	yes	$4 \times 10^{-4}$	0.9	0.999	0	BF16	256	4×4090
Qwen2.5-7B	20%	0.74	yes	$2 \times 10^{-4}$	0.9	0.999	0	BF16	512	8×A6000
	30%	0.61	yes	$2 \times 10^{-4}$	0.9	0.999	0	BF16	512	8×A6000

## A.2 MORE DETAILED EVALUATION RESULTS

To complement the summarized results presented in the main body, this section provides a more granular breakdown of our method’s performance on the LLaMA-3.2 model series. Table 10 and Table 11 present the detailed zero-shot evaluation scores for the LLaMA-3.2-1B and LLaMA-3.2-3B models, respectively. These tables show the performance across benchmarks, offering a comprehensive view that reinforces the consistent advantage of our entropy-based pruning criterion over existing baselines at both 20% and 30% sparsity ratios.

Table 10: Zero-shot performance of the compressed LLaMA-3.2-1B. The “**bolded**” represents the best result under the same pruning ratio.

Ratio	Method	ARCC	ARCE	BoolQ	Crows	OBQA	PIQA	Race	SiQA	TfQA	Wino	Average
0%	Llama-3.2-1.2B	37.1	60.7	64.0	62.6	37.6	74.2	37.6	42.9	37.7	60.4	51.5
20%	LLM-pruner	29.6	51.0	62.1	57.5	<b>33.6</b>	67.2	33.7	40.9	41.7	56.6	47.4
	Compresso	27.1	48.5	59.1	<b>58.6</b>	27.4	67.0	33.2	40.2	<b>43.7</b>	55.6	46.1
	LoRAPrune	29.6	51.0	66.2	54.4	29.6	66.4	34.1	42.0	41.7	<b>56.9</b>	47.2
	SDMPPrune	31.1	55.2	67.6	57.8	32.4	70.3	35.4	42.7	40.2	56.6	48.9
	HFPPrune(Ours)	<b>36.0</b>	<b>60.9</b>	<b>72.8</b>	55.4	33.0	<b>71.3</b>	<b>36.9</b>	<b>43.6</b>	41.3	56.6	<b>50.8</b>
30%	LLM-pruner	27.3	46.3	60.2	54.3	27.4	64.0	32.9	38.4	<b>42.5</b>	54.6	44.8
	Compresso	25.7	45.3	61.8	55.8	26.5	62.3	31.8	37.1	40.7	53.5	44.1
	LoRAPrune	28.1	46.4	59.6	56.3	27.6	64.9	32.4	37.3	41.6	54.4	44.9
	SDMPPrune	28.5	47.5	64.7	<b>56.9</b>	29.0	66.3	33.2	40.8	42.4	54.7	46.4
	HFPPrune(Ours)	<b>32.3</b>	<b>54.6</b>	<b>70.3</b>	55.9	<b>29.8</b>	<b>68.1</b>	<b>35.6</b>	<b>42.6</b>	40.4	<b>56.3</b>	<b>48.6</b>

Table 11: Zero-shot performance of the compressed llama-3.2-3B. The “**bolded**” represents the best result under the same pruning ratio.

Ratio	Method	BoolQ	Crows	OBQA	PIQA	Race	SiQA	TfQA	Average
0%	llama-3.2-3B	73.7	63.3	40.6	77.9	38.9	46.7	39.2	54.3
20%	LLMPruner	74.7	58.9	33.8	72.1	<b>40.1</b>	43.7	42.6	52.2
	LoRAPrune	68.8	57.9	33.1	71.4	35.3	42.9	42.6	50.3
	SDMPPrune	75.3	<b>59.3</b>	36.2	73.6	38.9	44.6	<b>45.8</b>	53.4
	HFPPrune(Ours)	<b>79.8</b>	58.8	<b>37.8</b>	<b>74.6</b>	39.0	<b>46.0</b>	42.6	<b>54.1</b>
30%	LLMPruner	66.2	54.9	32.8	66.2	37.1	41.3	42.3	48.7
	LoRAPrune	65.4	53.9	<b>33.8</b>	69.5	34.0	41.2	40.7	48.4
	SDMPPrune	72.8	<b>58.1</b>	30.4	70.7	36.8	43.3	<b>45.1</b>	51.0
	HFPPrune(Ours)	<b>77.6</b>	57.3	<b>33.8</b>	<b>71.0</b>	<b>38.6</b>	<b>46.5</b>	41.3	<b>52.3</b>

In addition to the main results which include a performance recovery stage, we provide a crucial study to assess the standalone effectiveness of our proposed pruning criterion. Table 12 presents the zero-shot performance of the Qwen series models immediately after pruning and without any fine-tuning. This retraining-free evaluation isolates the impact of the importance scoring itself. The results show that our entropy-based method consistently preserves higher accuracy than the baseline across various models and sparsity levels, validating our central hypothesis that our criterion is more effective at preserving a model’s intrinsic knowledge.

## A.3 QUALITATIVE ANALYSIS: DISTRIBUTION VISUALIZATION EXAMPLES

To visually illustrate the quantitative findings on distribution fidelity presented in the main text (Section 5.3.2), this section provides qualitative examples for four sample prompts. As shown in Figures 3, 4, 5, and 6, visualize and compare the next-token prediction distributions of three models: the original dense model, the model pruned with our Information Entropy (IE) criterion, and the model pruned with the traditional Cross-Entropy (CE) criterion. These examples visually confirm our quantitative results. A direct comparison shows that the distribution from the IE-pruned model

Table 12: Zero-shot performance of the compressed Qwen without any fine-tuning. The “**bolded**” represents the best result under the same pruning ratio.

Model Ratio	Method	ARCc	ARCe	BoolQ	Crows	OBQA	PIQA	Race	SiQA	TfQA	Wino	Average	
Qwen2.5-7B	0% Dense	51.1	77.5	84.7	64.0	46.8	79.8	41.4	54.8	56.4	73.2	63.0	
	20%	SDMPrune	37.5	58.5	<b>74.3</b>	61.2	38.7	76.1	37.1	<b>50.4</b>	<b>50.2</b>	<b>63.5</b>	<b>54.8</b>
		HFPrune(Ours)	<b>39.9</b>	<b>66.9</b>	68.2	<b>61.8</b>	<b>39.8</b>	<b>77.5</b>	<b>37.2</b>	48.2	46.1	60.1	54.6
	30%	SDMPrune	30.0	52.8	<b>66.4</b>	<b>59.3</b>	32.8	71.2	33.6	42.6	<b>44.9</b>	<b>58.9</b>	49.2
		HFPrune(Ours)	<b>34.0</b>	<b>58.3</b>	61.9	58.3	<b>35.2</b>	<b>74.1</b>	<b>35.2</b>	<b>44.3</b>	43.1	56.2	<b>50.1</b>
	Qwen2.5-1.5B	0% Dense	45.4	72.0	73.0	60.9	40.4	76.0	36.5	48.9	46.6	63.5	56.3
20%		SDMPrune	30.6	55.6	<b>65.5</b>	54.5	32.2	70.2	32.9	41.0	<b>43.5</b>	56.7	48.3
		HFPrune(Ours)	<b>32.3</b>	<b>56.8</b>	64.0	<b>56.1</b>	<b>34.0</b>	<b>71.9</b>	<b>34.7</b>	<b>42.6</b>	41.1	<b>57.0</b>	<b>49.1</b>
30%		SDMPrune	23.2	40.5	57.9	<b>54.4</b>	28.8	65.6	32.3	38.5	<b>44.2</b>	<b>54.8</b>	44.0
		HFPrune(Ours)	<b>25.9</b>	<b>46.8</b>	<b>62.0</b>	50.8	<b>31.0</b>	<b>69.0</b>	32.3	<b>39.9</b>	42.1	53.7	<b>45.3</b>
Qwen3-1.7B		0% Dense	42.8	70.0	77.6	55.2	37.8	71.9	36.8	44.9	46.0	61.6	54.4
	20%	SDMPrune	23.3	36.5	51.0	49.6	26.4	63.2	32.1	36.6	<b>43.9</b>	54.1	41.7
		HFPrune(Ours)	<b>27.8</b>	<b>44.2</b>	<b>64.6</b>	<b>51.8</b>	<b>28.6</b>	<b>65.2</b>	<b>34.4</b>	<b>38.6</b>	42.0	<b>54.5</b>	<b>45.2</b>
	30%	SDMPrune	<b>21.0</b>	32.1	39.1	47.9	<b>25.6</b>	58.5	<b>28.7</b>	35.2	<b>46.0</b>	<b>51.1</b>	38.5
		HFPrune(Ours)	20.0	<b>35.3</b>	<b>62.1</b>	47.9	25.2	<b>59.7</b>	27.1	<b>35.5</b>	43.8	50.5	<b>40.7</b>

more closely mirrors the original, both in its overall shape (corresponding to lower JS Distance) and in its preservation of high probability tokens (corresponding to higher Jaccard Similarity). This demonstrates the superior ability of our IE criterion to maintain the model’s predictive fidelity.

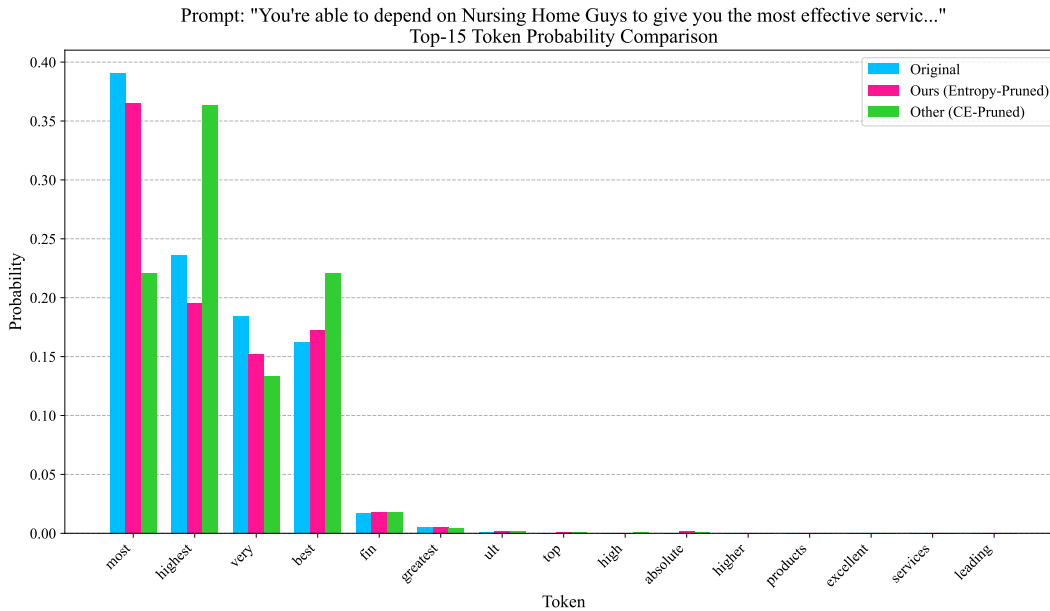


Figure 3: Sample1, You are able to depend on Local Bathroom Remodel Crew to deliver the very best expert services when it comes to Bathroom Remodeling in Williamstown, NJ. Our crew of experienced experts will provide the expert services that you require with the most innovative technologies around. We make sure

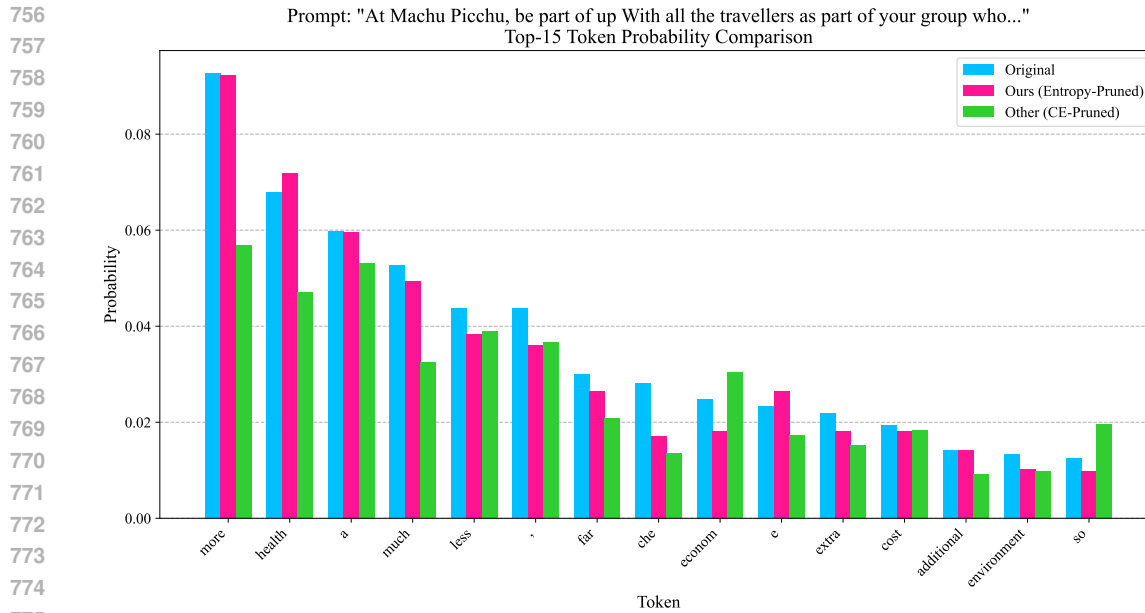


Figure 4: Sample2, At Machu Picchu, be part of up With all the travellers as part of your group who hiked the common Inca Trail. If skies are very clear, get pleasure from a spectacular views more than The traditional metropolis from your Sunlight Gate, right before going on a guided stroll round the ruins. Mobile Massage Can Save The Environment! A more pure, not to say

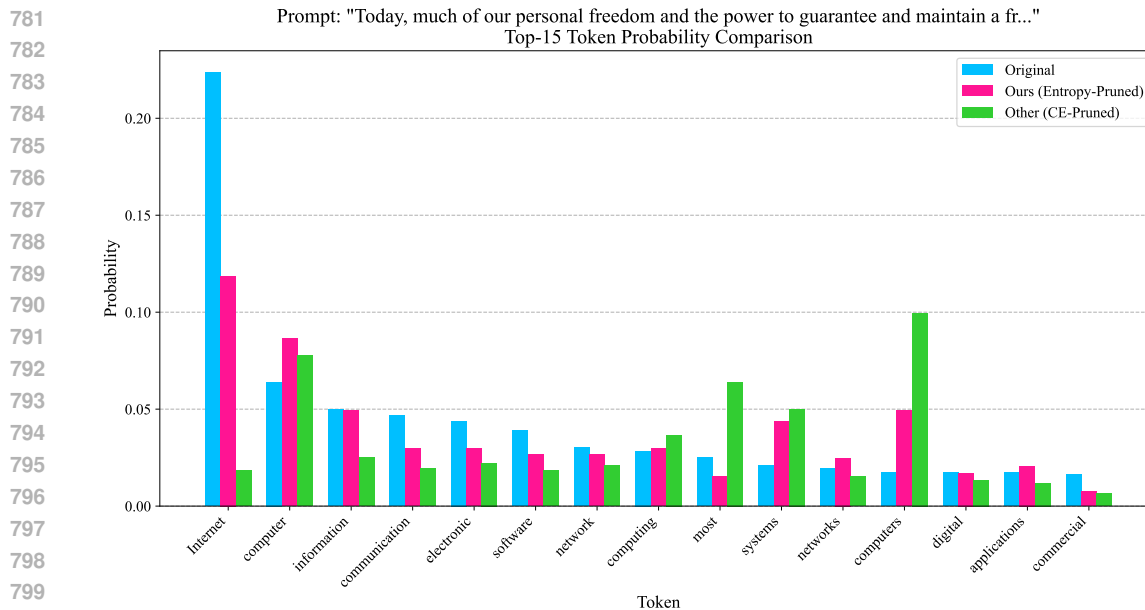


Figure 5: Sample3, Today, much of our personal freedom and the power to guarantee and maintain a free society depends on cryptographic primitives (e.g., digital signatures and encryption) incorporated in the security protocols of today's

#### A.4 LLM USAGE

We utilized Large Language Model to refine the phrasing of this paper. The core contributions of this work including the motivation, core idea, method design, and experimental design are entirely the work of the authors.

810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
820  
821  
822  
823  
824  
825  
826  
827  
828  
829  
830  
831  
832  
833  
834  
835  
836  
837  
838  
839  
840  
841  
842  
843  
844  
845  
846  
847  
848  
849  
850  
851  
852  
853  
854  
855  
856  
857  
858  
859  
860  
861  
862  
863

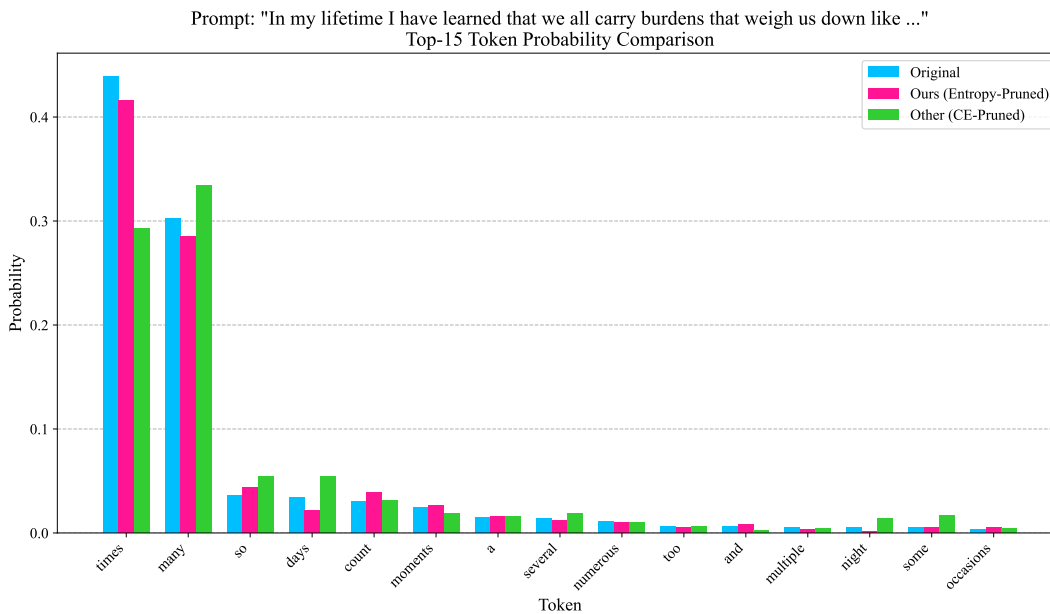


Figure 6: Sample4, In my lifetime I have learned that we all carry burdens that weigh us down like rocks threatening to squash us like bugs. I have also learned that the only true way to feel that burden lightened, to feel strong and capable of enduring is through the Lord Jesus Christ’s Atonement. He suffered not only for our sins but for our physical pains, our doubts, our fears and regrets. Every burden we carry he has felt and taken upon his back. There have been