

---

# DarwinLM: Evolutionary Structured Pruning of Large Language Models

---

Shengkun Tang<sup>1</sup> Oliver Sieberling<sup>2</sup> Eldar Kurtic<sup>3,4</sup> Zhiqiang Shen<sup>1</sup> Dan Alistarh<sup>3,4</sup>

## Abstract

Large Language Models (LLMs) have achieved significant success across various NLP tasks. However, their massive computational costs limit their widespread use, particularly in real-time applications. Structured pruning offers an effective solution by compressing models and directly providing end-to-end speed improvements, regardless of the hardware environment. Meanwhile, different components of the model exhibit varying sensitivities towards pruning, calling for *non-uniform* model compression. However, a pruning method should not only identify a capable sub-structure, but also account for post-compression training. To this end, we propose *DarwinLM*, a method for *training-aware* structured pruning. *DarwinLM* builds upon an evolutionary search process, generating multiple offspring models in each generation through mutation, and selecting the fittest for survival. To assess the effect of post-training, we incorporate a lightweight, multistep training process within the offspring population, progressively increasing the number of tokens and eliminating poorly performing models in each selection stage. We validate our method through extensive experiments on Llama-2-7B, Llama-3.1-8B and Qwen-2.5-14B-Instruct, achieving state-of-the-art performance for structured pruning. For instance, *DarwinLM* surpasses ShearedLlama while requiring  $5\times$  less training data during post-compression training. Code and all weights are released at: <https://github.com/IST-DASLab/DarwinLM>.

## 1. Introduction

The high accuracy of Transformer-based models on a wide range of tasks comes with massive computational requirements, which hinders deployability. A long line of research has been conducted to reduce the computational cost of

large language models via methods such as quantization (Frantar et al., 2022; Dettmers et al., 2023), pruning (Xia et al., 2024; Frantar & Alistarh, 2023) and distillation (Hsieh et al., 2023). In this work, we explore structured pruning on large language models (LLMs) by removing whole rows or columns in a weight matrix, resulting in regular but smaller matrices. Unlike unstructured pruning (Frantar & Alistarh, 2023), the model produced by structured pruning can be accelerated on mainstream hardware without any specific design for computation.

While conventional pruning methods generally prune each layer or block uniformly, *non-uniform compression methods*, e.g. (Yin et al., 2023; Sieberling et al., 2024), showed that different LLM layers can have significantly different “sensitivities” to pruning, which can be leveraged to obtain higher compression while preserving accuracy. To address this, smaller-scale pruning methods such as ZipLM (Kurtić et al., 2024) propose to utilize a dynamic-programming-based search (Frantar & Alistarh, 2022) to obtain a sparse model with runtime guarantees. However, there are several challenges for methods such as ZipLM on large-scale models such as Llama-2-7B (Touvron et al., 2023): for instance, ZipLM only considers the local layer-wise error during the search, which is not consistent with performance on in-context learning (ICL) or downstream tasks.

**Overview.** In this paper, we propose a new structured pruning method based on evolutionary search called *DarwinLM*, which addresses these issues. Specifically, the search starts from a “parent” model, generated by pruning the original model using second-order information. In each search step, it then generates “offspring” candidate models by copying the parent and “shifting” sparsity from one layer to another, by what we call a *level switch mutation*. Moreover, to obtain the optimal sparse model given additional training, we use a small-scale dataset to finetune each generated offspring, and select the best offspring after finetuning. Empirically, our method scales to LLMs, and achieves state-of-the-art performance on both one-shot and post-training settings compared with previous pruning methods. For instance, we obtain smaller and better sparse models on both Llama-2, Llama-3.1 and Qwen-2.5 compared with the state-of-the-art ZipLM pruning method; furthermore, on Llama-2-7B, our sparse model with 2.7B parameters outperforms the ShearedLlama-2.7B structured pruned model, while using

<sup>1</sup>Department of Machine Learning, MBZUAI, Abu Dhabi, UAE <sup>2</sup>ETH Zurich, Zurich, Switzerland <sup>3</sup>ISTA, Vienna, Austria <sup>4</sup>Red Hat AI, Boston, USA. Correspondence to: Dan Alistarh <dan.alistarh@ist.ac.at>.

---

**DarwinLM: Evolutionary Structured Pruning of Large Language Models**


---

$5\times$  less data for post-compression training (50B Tokens vs. 10B Tokens). Moreover, a *DarwinLM* with 4.6B parameters pruned from Llama-3.1 surpasses an OLMO model with 7B parameters pre-trained on  $250\times$  more data (10B Tokens vs 250T tokens). We also compare our method with a line of more coarse-grained structured pruning methods including ShortGPT (Men et al., 2024), Shortened-Llama (Kim et al., 2024), and EvoPress (Sieberling et al., 2024) in a one-shot setting, showing that *DarwinLM* provides better performance across compression rates.

We summarize our contributions as follows:

- We propose a new structured pruning method that unifies second-order pruning and evolutionary search, with the goal of producing the “optimal” non-uniformly “thinned” model given a fixed size or runtime constraint. Importantly, the resulting models can be run efficiently on any hardware, and can meet hardware-specific performance constraints.
- At the technical level, we introduce a pruning technique that leverages second-order information with a training-aware offspring selection technique, to identify the non-uniform sparse model that meets the performance requirement and also has the highest potential for accuracy gains during post-compression training. Empirical results show that this approach enhances performance both in terms of loss, but also in terms of performance on downstream tasks.
- We validate the effectiveness of *DarwinLM* through extensive experiments. For instance, our method compresses Llama-2-7B to 2.7B parameters, Llama-3.1-8B to 4.6B parameters, and Qwen-2.5-14B-Instruct to 8.4B with much lower performance degradation compared to prior state-of-the-art methods. For instance, our approach outperforms the state-of-the-art ShearedLlama method for structured pruning, while using  $5\times$  fewer data tokens for post-training. Importantly, our method can produce accurate, tailored compressed models: for example, we obtain a 4.6B-parameter variant of Llama-3.1-8B that significantly outperforms the 7B-parameter OLMO model trained on  $250\times$  more data in terms of accuracy on standard zero-shot tasks.

## 2. Related Work

**Structured Pruning Methods.** Structured pruning methods for LLMs (Wang et al., 2020; Ma et al., 2023; Xia et al., 2024; Kurtić et al., 2024; Men et al., 2024; Kim et al., 2024; An et al., 2024; Muralidharan et al., 2024) typically focus on pruning along the depth dimension (dropping layers) or on pruning width (such as attention heads, and MLP intermediate dimensions), and in some cases, both depth and width. Among recent advances, the state-of-the-art is provided by ShearedLLaMA (Xia et al., 2024), which utilizes targeted structured pruning, which reduces a larger model to

a specified target shape by pruning layers, attention heads, and intermediate or hidden dimensions in an end-to-end process that is split into regularized fine-tuning, pruning, and further fine-tuning. In addition, it implements *dynamic batching*, which adjusts the composition of sampled data in each training batch dynamically, based on varying loss proxies across different evaluation domains. By contrast, *DarwinLM* provides more accurate structured pruning, combining evolutionary search and second-order information, requiring only a fraction of the data to recover accuracy.

The recent work on MINITRON (Muralidharan et al., 2024) established a set of effective compression practices for LLMs by integrating depth, width, attention, and MLP pruning with knowledge distillation (KD)-based retraining. These practices are derived from an in-depth empirical exploration of pruning strategies across each axis, methods for combining different axes, distillation approaches, and search techniques to identify optimal compressed architectures. Our contributions are orthogonal to MINITRON, as we mainly investigate more accurate pruning techniques—many of their findings should also transfer to our setting. Unfortunately, MINITRON uses a closed fine-tuning dataset, and did not release pre-fine-tuning checkpoints. Thus, we are unable to compare pruning techniques and end models relative to MINITRON.

**Non-uniform Pruning Methods.** The distribution of importance across depth, attention heads, and width in the model varies between layers and is not uniform. Low-importance modules tend to be concentrated in specific locations and positions within the model. In the LLM domain, Klein et al. (2023) utilized multi-objective NAS to compress LLMs while optimizing their performance for fine-tuning downstream tasks. SIMPLE (Tao et al., 2023) detects redundant network structures by applying learnable masks to all compressible components, followed by sparse training. Unlike these prior methods, we propose a novel evolutionary framework, that maintains a fixed sparsity across the search process and integrates the effectiveness of finetuning within the fitness evaluation. Therefore, *DarwinLM* is *training-aware*. From the technical standpoint, our work builds upon the recent EvoPress approach (Sieberling et al., 2024). EvoPress was designed for *non-uniform unstructured pruning*, *non-uniform quantization*, and *layer dropping*, with a focus on achieving a target model size in a *one-shot setting*. By contrast, *DarwinLM* differs by focusing on fine-grained structured pruning (at the level of rows/columns), and incorporating the effect of continued training into the compression allocation. The more fine-grained structured pruning we employ significantly improves performance, while maintaining guaranteed speedup without specific hardware support (contrary to e.g. unstructured sparsity). Additionally, two equally performing pruned models can respond differently to continued training, which motivates integrating a lightweight finetune into the search process.

## DarwinLM: Evolutionary Structured Pruning of Large Language Models

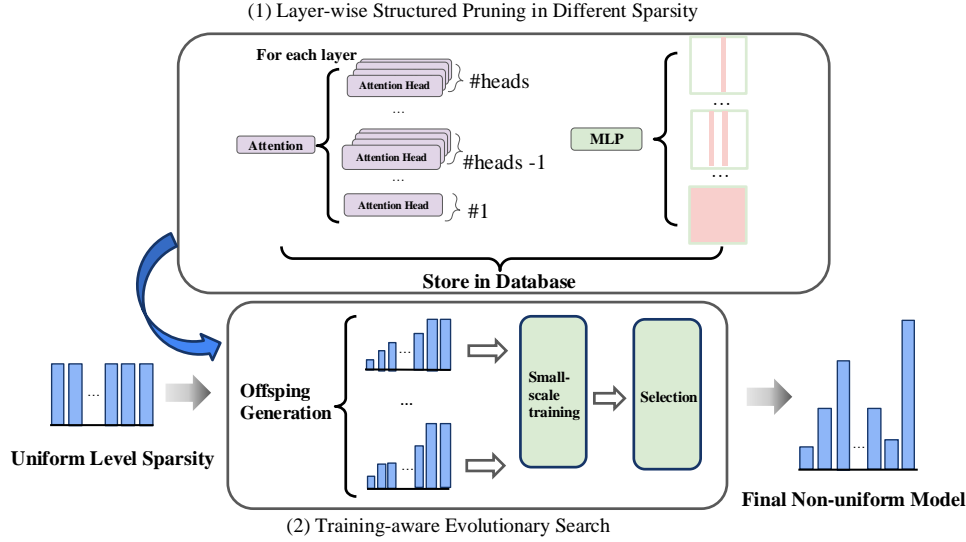


Figure 1. Visual illustration of *DarwinLM* pipeline. 1) generate sparsity level database with different sparsities by second-order structured pruning. 2) evolutionary search with training-aware selection based on the sparsity level database.

**Other Compression Methods.** Several approaches have been explored in the literature to reduce computational and memory requirements of LLMs without significantly degrading performance, including knowledge distillation, quantization, binarization, and sparsity. In knowledge distillation (Hinton et al., 2015; Sanh, 2019; Gu et al., 2024; Liu et al., 2024; Xu et al., 2024a), a smaller, simpler model (the “student”) is trained to replicate the behavior of a larger, more complex model (the “teacher”). The goal is to transfer the knowledge from the teacher to the student while retaining most of the performance benefits of the larger model. Quantization (Xiao et al., 2023; Lin et al., 2024; Li et al., 2024; Wang et al., 2023; Huang et al., 2024; Xu et al., 2024b; Ma et al., 2024; Tang et al., 2024) reduces the precision of model weights and activations. While this can dramatically reduce the model size and computation, the challenge lies in maintaining accuracy.

### 3. Method

We begin with outlining the problem formulation in Section 3.1, followed by a description of the pruning method in Section 3.2. We discuss the generation of the sparsity level database, which is used for the search, in Section 3.3. We describe the fitness environment in Section 3.4 and the evolutionary algorithm in Section 3.5. An overview of the pipeline is provided in Figure 1.

#### 3.1. Problem Setup

Given a target sparsity, *DarwinLM* aims to find the model with the best sparsity allocation adhering to this constraint. Formally, let  $S$  denote the target sparsity factor, let  $M_{base}$  be the base model. Then, our problem is reduced to:

$$\hat{M} = \arg \max_M f(M) \quad (1)$$

under the constraint  $S(M) \geq S$ , where  $f(\cdot)$  is a function that evaluates the performance of a model and  $S(\cdot)$  measures the sparsity of a given model. However, Equation 1 presents a non-differentiable optimization problem. To address this, we will employ an evolutionary algorithm to search for the optimal sparse structure.

#### 3.2. Second-Order Structured Pruning

We follow ZipLM (Kurtić et al., 2024), which itself is an extension of Optimal Brain Surgeon (OBS) (Hassibi & Stork, 1992; Kurtic et al., 2022), for *layer-wise structured pruning* of a pre-trained LLM. Specifically, for each layer, given a calibration dataset  $X$  of layer inputs and the original layer weights  $W$ , we aim to find

$$\arg \min_{\hat{W}} \|WX - \hat{W}_{:,M}X\|_2 \quad (2)$$

subject to  $\hat{W}_{:,M} \in \mathcal{C}$ , where  $M$  refers to a *column mask* and  $\mathcal{C}$  is the compression constraint. To ensure that the sparse weights  $\hat{W}$  produce outputs similar to those of the original weights  $W$ , we must not only identify the less significant structures for pruning, but also compute an update  $\delta$  for the remaining weights to compensate for the error introduced by pruning. For this purpose, denote by  $H = XX^T$  the Hessian matrix for the  $\ell_2$ -minimization problem in Equation 2. Define  $W_{i,M}$  as the weight in row  $i$  masked by  $M$  and let  $(H^{-1})_{M,M}$  be the submatrix of the inverse Hessian corresponding to the entries under the mask  $M$ . Now, we can compute the optimal structured mask with corresponding weight updates  $\delta$  by:

$$\arg \min_M \sum_{i=1}^{d_{row}} W_{i,M} \cdot (H_{M,M}^{-1})^{-1} \cdot W_{i,M}^T \quad (3)$$

$$\delta = -W_{:,M} \cdot (H_{M,M}^{-1})^{-1} \cdot H^{-1}_{M,:} \quad (4)$$

---

**DarwinLM: Evolutionary Structured Pruning of Large Language Models**


---

This formulation extends the derivation of OBS to account for all rows  $d_{row}$ . In our context, we focus on three types of pruned structures: (1) head pruning in multi-head self-attention, (2) pruning of the intermediate dimension of MLP modules, and (3) removing the entire attention or MLP module.

**Granularity.** To reduce the required memory for storing the database, we enforce the number of pruned dimensions in the MLP modules to be a multiple of  $m = 32$ . For attention modules, we prune on a per-head basis. For each module, we only consider identifying the pruned columns of the final output matrix, referring to the down projection in the case of an MLP. Once the pruned structure of the output matrix is determined, the corresponding rows are pruned in the other matrices (i.e., the K, Q, and V matrices in the attention module, and the up and gate projections in the MLP). However, if the model applies group-query attention (GQA) (Ainslie et al., 2023), such as in Llama-3.1 and Qwen-2.5, we avoid pruning the K and V matrices. During the forward pass, we remove the corresponding heads in the repeated K and V matrices to obtain computationally compatible structures and reduce computation.

### 3.3. Sparsity Level Database

We first generate a sparsity level database, which will be used in the evolutionary search, where candidate models are stitched together from their respective unitwise sparsity levels. However, we must ensure that all considered models in this search process adhere to the targeted sparsity factor. We will do so by initializing the search with a valid model and then applying a sparsity-preserving mutation operator. For this purpose, we generate the database such that the difference in terms of (absolute) sparsity between two levels is consistent across all levels and modules. (In our implementation, all attention / MLPs employ the same sparsity step size, but the step size for attention differs from that of MLPs.) Thus, we can mutate a model while maintaining the targeted sparsity by simply increasing the same number of levels as we decrease.

Suppose we aim to obtain a database with  $N_l$  sparsity levels, then the number of pruned heads in attention modules and the pruned rows in MLP modules are as following:

$$\text{Round}\left(\frac{i \times N_{head}}{N_l}\right), \quad (5)$$

$$m \times \text{Round}\left(\frac{i \times N_{inter}}{N_l}\right), i \in 0, 1 \dots N_l \quad (6)$$

where  $N_{head}$ ,  $N_{inter}$  refer to the number of heads in the attention module and the intermediate size in the MLP module, while  $m$  is the pruning granularity of the MLP module. As a result, we can obtain the sparsity level database as:

$$D = \{i : \mathbf{W}_i\}, i \in 0, 1 \dots N_l \quad (7)$$

---

**Algorithm 1** DarwinLM: Evolutionary Search with training-aware offspring selection.

---

**Input:**

$N$ : number of generation.  $S$ : selection steps.

$\lambda$ : number of offsprings in each generation.

$T_f$ : tokens for finetuning.  $T_s$ : tokens for selection.

**Initialization:**

$D \leftarrow \text{databaseGen}()$

## Sampled levels are all intergers

$parent \leftarrow \text{UniformLevelSample}()$

**Optimization:**

## Offspring Generation via Mutation

**for**  $t \leftarrow 1$  to  $N$  **do**

$candidates \leftarrow [parent]$

**for**  $i \leftarrow 1$  to  $\lambda$  **do**

$offspring \leftarrow \text{LevelSwitchMutation}(parent)$

$candidates.append(offspring)$

**end for**

**end for**

    ## Selection:

**for**  $Step \leftarrow 1$  to  $S$

$cand\_models = []$

**for**  $candidate \in candidates$  :

$cand\_model \leftarrow \text{stitch}(candidate, D)$

$cand\_model \leftarrow \text{train}(cand\_model, T_f[step])$

$cand\_models.append(cand\_model)$

**end for**

$candidates \leftarrow \text{selectTopKFit}(cand\_models, T_s[step])$

**end for**

**return**  $candidates[0]$

---

The database will be used to stitch the model given the sparsity level for each module.

### 3.4. Fitness Environment

Finding a suitable fitness environment is fundamental to evolutionary search. Although models are typically evaluated based on their performance on downstream tasks, this approach is impractical in our context due to the lengthy evaluation times and the risk of overfitting. As an alternative, we follow EvoPress (Sieberling et al., 2024) and adapt the Kullback-Leibler (KL) divergence between the outputs of the dense model and sparse model on a small-size calibration data as a metric to evaluate the fitness of a candidate. Given a sparsity level database as described in the previous section, every possible model in our search space is defined by its module-wise sparsity levels. Therefore, by assuming  $\mathbf{M}$  to be an  $n$ -tuple over the set of sparsity levels  $[N_l]$ , we can rewrite our objective function as

$$\hat{\mathbf{M}} = \arg \min_{\mathbf{M}} D_{KL}(\mathbf{M}) \quad (8)$$

subject to  $\mathcal{S}(\mathbf{M}) \geq \mathcal{S}$ .

## DarwinLM: Evolutionary Structured Pruning of Large Language Models

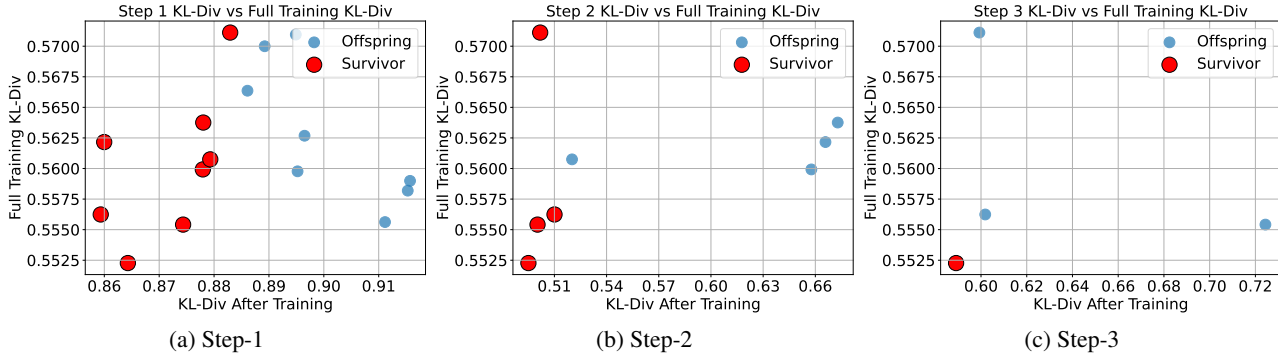


Figure 2. Motivation of the training-aware selection. The Y-axis depicts the KL-Divergence of the model after full post-training while x-axis is the KL-Divergence after small-scale data training. The results indicate that our training-aware selection can select the best offspring for large-scale training.

### 3.5. Evolutionary Search Algorithm

We approach the problem in Equation (1) via evolutionary search. The algorithm, described in Algorithm 1, includes initialization, mutation process, and multi-step training-aware offspring selection process.

**Initialization.** We first generate the sparsity level database composed of pruned layers for different sparsities as discussed in Sections 3.2 and 3.3. We perform the level increase and decrease, the mutation, separately in the attention and MLP modules. Initially, our search algorithm starts from uniform compression. For example, if the target sparsity level is 5, then there are two tuples to record the level for each module:  $[5 \cdots 5]_{attn}$  and  $[5 \cdots 5]_{mlp}$ . In the case of gradual pruning, we compute the residual value between the target sparsity level in different stage and randomly add the residual value to the searched results from the previous stage.

**Mutation Process.** We apply the mutation operator within each module group, meaning, we never exchange sparsity levels between an attention and a MLP module. Given a fixed step size, the operator randomly picks one unit to increase the sparsity level with the step size. After that, another random unit is selected to decrease the sparsity level with the same step size. Based on the database design, we can guarantee that all generated offspring meet the requirement of the compression constraint.

**Multi-step Training-aware Selection Process.** Our goal is not only to find the best sparse model in a one-shot setting, but to account for continued training. We start from the observation that training on little data is a good predictor of more large-scale training. We demonstrate this in Figure 2, where we generate 16 offspring for Llama2-7B. We first use 2M tokens to train all offspring as a “large-scale” full training. We apply 3 selection steps, each with [8, 4, 1] survivors respectively. We utilize [1024, 2048, 8196] tokens for selection and [10K, 50K, 200K] tokens for training. As

depicted in Figure 2 (right), the best offspring after full finetuning is successfully identified in the selection process. This motivates what we call *training-aware offspring selection*, meaning that we incorporate a lightweight finetune into the selection process. Given a small-scale training data  $T_f$  (less than 200K tokens in our experiments), we train each offspring and compute the fitness value of models after training. To ensure a robust and efficient training and selection, we employ a multi-step training and selection process, following the approach outlined in prior work (Sieberling et al., 2024). Specifically, the training and selection are performed iteratively over  $S$  rounds. In each round, a progressively smaller subset of offspring is retained, while the number of samples for fitness evaluation and training is increased for each step. The final surviving candidate is selected as the starting point for the next generation.

## 4. Experiments

### 4.1. Setup

**Models and Datasets.** Given a target sparsity level and a set of pre-trained weights, our method searches for combinations of per-layer sparsity levels under the sparsity constraint, based on a small generic calibration set. In our experiments, we consider Llama-2-7B (Touvron et al., 2023), Llama-3.1-8B (Dubey et al., 2024) and Qwen-2.5-14B-Instruct. We utilize the open-source dataset Fineweb-Edu (Lozhkov et al., 2024) for both calibration and post-training. The dataset is filtered according to the sample score provided with the dataset. All samples with a lower score than 0.9 are removed from the dataset, resulting in a dataset with 80B tokens. For the search process, we use at most 16 sequences for calibration, making this process highly lightweight. The finetuning data for the offspring models is also sampled from the Fineweb-Edu dataset.

**Baselines.** First, we compare our non-uniform sparse model with a uniform sparse model under a similar computational

## DarwinLM: Evolutionary Structured Pruning of Large Language Models

Table 1. Comparison of main results for *DarwinLM* and baseline methods on LLaMA2-7B. Our method achieves the best average performance on benchmarks compared to the baseline methods. With only 10B tokens of fine-tuning, our method beats ShearedLlama, which is fine-tuned with 50B tokens. (†) refers to training on the same 10B tokens we use. We omit MMLU results in this table as they are close to random for all Llama2 models.

Method	Param.	SciQ	PIQA	WG	ArcE	ArcC	HS	LogiQA	BoolQ	Avg
Dense	6.7B	93.7	78.1	69.3	76.4	53.0	78.6	30.7	77.7	69.2
Uniform	3.4B	44.1	57.1	53.3	33.5	32.2	27.3	25.0	49.0	40.1
ZipLM	4.0B	87.4	64.4	<b>58.3</b>	53.2	33.6	50.1	25.5	63.6	54.5
ShearedLLaMA	2.7B	84.5	66.4	53.4	49.8	28.4	47.6	27.6	50.9	51.0
<i>DarwinLM</i> (one-shot)	2.7B	<b>85.6</b>	<b>70.8</b>	55.8	<b>63.3</b>	<b>38.1</b>	<b>53.2</b>	<b>28.5</b>	<b>62.7</b>	<b>57.2</b>
ShearedLLaMA (50B)	2.7B	90.8	75.8	64.2	67.0	41.2	70.8	28.2	63.0	62.6
ShearedLLaMA (10B†)	2.7B	92.0	73.6	63.1	69.8	42.0	64.4	29.0	62.1	61.9
<i>DarwinLM</i> (10B)	2.6B	<b>90.8</b>	72.2	<b>65.1</b>	68.5	<b>45.0</b>	67.2	28.5	<b>64.6</b>	<b>62.8</b>

budget. Additionally, on Llama-2-7B, we conduct comparisons with ZipLM (Kurtić et al., 2024) and ShearedLlama (Xia et al., 2024). ZipLM employs dynamic programming to search for the sparse model structure, while ShearedLlama learns pruning masks for Llama-2-7B’s weights and applies large-scale fine-tuning on 50B tokens. We perform the evaluation using the publicly available weights after pruning and fine-tuning, as provided by the respective papers. For ZipLM, we reproduce their implementation at a larger scale, following the original paper’s methodology. We limit our comparison with ShearedLlama to Llama-2-7B, as the original paper only reports results for this model, and the tuning costs for adapting it to other models are substantial. We also compare *DarwinLM* in a one-shot setting against other one-shot structured pruning methods, including EvoPress (Sieberling et al., 2024), ShortGPT (Men et al., 2024), Shortened Llama (Kim et al., 2024), and a block scoring approach using sliding window cosine similarity (Gromov et al., 2024). All of these methods perform structured pruning on a per-module or per-layer level. Finally, we also add OLMO (Groeneveld et al., 2024) series model, which is the open-source large-scale pre-trained dense model, as the reference to compare our post-compression training model. Specifically, OLMO-7B model is trained with 2.5T tokens while OLMO-0424-7B utilizes 2.05T tokens for training. We use the official pre-trained weight released in Huggingface for evaluation.

**Evaluation.** We follow ShearedLlama (Xia et al., 2024) to evaluate our method on several downstream tasks including 0-shot accuracy on ARC-easy (Clark et al., 2018), LogiQA (Liu et al., 2020), PIQA (Bisk et al., 2020), SciQ (Welbl et al., 2017), BoolQ (Clark et al., 2019), 5-shot on MMLU (Hendrycks et al., 2020) and WinoGrande (Sakaguchi et al., 2021), 10-shot on HellaSwag (Zellers et al., 2019) and 25-shot on ARC Challenge (Clark et al., 2018). We utilize the *lm-evaluation-harness* framework (Gao et al.) to evaluate all downstream tasks.

**Implementation Details.** When generating the sparsity level database, we set the minimum and maximum levels are 0 and 10, which indicates the 0% and 100% sparsity.

On Llama-2-7B, we first prune the model with a target sparsity level 5 in the one-shot setting using 2048 calibration samples and fine-tune the sparse model on 10B tokens. After that, we continue to prune the model to target sparsity level 6 based on the fine-tuned model with 2K calibration data. We prune Llama-3.1-8B and Qwen-2.5-14B-Instruct models with target sparsity level 5. The final pruned models are trained on an additional 10B Fineweb tokens. For the evolutionary search, we set the number of generation to 200. For each generation, we generate  $\lambda = 16$  offspring for selection. During selection, we apply 4-steps selection with [1024, 2048, 4096, 8192] token for fitness computation and [10K, 50K, 100K, 200K] token for offspring finetuning. The learning rate of training during the searching is  $1e-5$ . The pruning and search process is conducted on a  $10\times$  L40s GPU workstation. Our training code is based on the LLM-Foundry repository. Our batch size is 1,024 for Llama-2, 1152 for Llama-3.1, and 2048 for Qwen-2.5. The base learning rate is  $1e-4$  with a cosine decay strategy. The post-training with 10B tokens is operated on a  $40\times$  H100 cluster for 13 hours. More details of the pruning and training can be found in the Appendix.

## 4.2. Main Results

**Results on Llama-2-7B.** We prune the Llama-2-7B model down to 2.7B with a target level 6. The main results are shown in Table 1. For the pruned models, our method achieves the highest performance on all downstream tasks, except for WinoGrande, where the ZipLM includes much more parameters. Our method also attains the highest average score. In contrast, the uniform pruning method results in a significant performance drop, with an average accuracy of only 40.1, essentially a performance collapse compared to the dense model. Specifically, the uniformly-pruned model generates nearly random results on benchmarks such as WinoGrande, HellaSwag, LogiQA, BoolQ, and MMLU. By contrast, *DarwinLM* achieves an average score of 57.2, outperforming ZipLM (54.5 with 4.0B parameters) and ShearedLlama (51.0 with 2.7B parameters). This comparison highlights the effectiveness of non-uniform

## DarwinLM: Evolutionary Structured Pruning of Large Language Models

Table 2. Comparison of results for *DarwinLM* and baseline models on Llama-3.1-8B. With similar speedup, our method achieves the best performance on all benchmarks compared with uniform pruning and ZipLM. After post-training with 10B tokens, the performance recovers from 54.0 to 58.1.

Model	Method	Param.	SciQ	PIQA	WG	ArcE	ArcC	HS	LogiQA	BoolQ	MMLU	Avg
Llama-3.1-8B	Dense	8B	96.3	81.2	74.3	81.4	58.2	81.7	31.1	84.0	65.2	72.8
	Uniform	4.5B	29.1	53.6	51.7	26.0	23.6	27.1	25.5	62.1	25.7	36.1
	ZipLM	6B	65.5	60.6	56.0	40.2	36.2	34.4	28.1	63.0	27.9	45.7
	<i>DarwinLM</i> (one-shot)	4.6B	<b>84.9</b>	<b>69.4</b>	<b>57.3</b>	<b>59.6</b>	34.2	<b>44.6</b>	24.1	62.2	<b>28.5</b>	<b>51.6</b>
	OLMO (2.5T)	7B	92.8	79.4	70.4	73.3	44.9	77.1	27.9	72.5	28.3	62.9
<i>DarwinLM</i> (10.0B)	4.6B	<b>93.2</b>	74.8	67.4	73.2	<b>51.6</b>	71.3	<b>30.7</b>	71.1	<b>40.6</b>	<b>63.7</b>	
Qwen-2.5-14B-Instruct	Dense	14B	96.8	81.9	79.1	85.7	72.8	85.1	38.5	87.9	80.0	78.6
	Uniform	8.6B	78.2	72.7	57.6	76.1	45.6	47.0	28.1	61.6	45.5	56.9
	ZipLM	8.5B	69.0	66.4	52.8	60.1	38.3	43.3	29.6	60.2	25.0	49.4
	<i>DarwinLM</i> (one-shot)	8.4B	<b>84.3</b>	<b>73.9</b>	<b>60.5</b>	75.7	<b>48.0</b>	<b>53.3</b>	29.3	<b>66.9</b>	<b>43.1</b>	<b>59.4</b>
	OLMO-0424 (2.05T)	7B	96.1	80.1	72.1	73.8	49.2	78	29.3	80.8	52.1	67.9
<i>DarwinLM</i> (10.0B)	8.4B	89.5	78.1	70.7	<b>79.6</b>	<b>57.6</b>	74.9	<b>33.5</b>	73.9	<b>57.9</b>	<b>68.4</b>	




Figure 3. Performance comparison of *DarwinLM* and ShearedLlama with different training token numbers. *DarwinLM* achieves better performance than ShearedLlama on all training token number settings.

structured pruning, particularly at high sparsity. After post-compression training, the pruned models see a significant recovery in performance. Notably, with only 10B tokens for training, *DarwinLM* reaches an average score of 62.8, surpassing the 62.6 reported by ShearedLlama, which was trained with 50B tokens. Furthermore, when we train the pruned model released by ShearedLlama under the same conditions and with 10B tokens, it achieves an average score of 61.9, which is considerably lower than *DarwinLM*. Figure 3 presents the training performance curve of *DarwinLM* and ShearedLlama, showing the average performance on benchmarks evaluated every 2B tokens. More detailed results for individual tasks can be found in the Appendix. Overall, with the same amount of training data, *DarwinLM* consistently outperforms ShearedLlama. Interestingly, with just 2B tokens, *DarwinLM* produces comparable results to ShearedLlama trained on 10B tokens. The curve demonstrates that *DarwinLM* achieves superior performance to ShearedLlama regardless of the number of training tokens.

**Results on Llama-3.1-8B.** We also pruned the Llama-3.1 8B model to 4.6B parameters with a target sparsity level 5. The comparison results are shown in Table 2. Similar to

Table 3. One-shot result comparison of the uniformly pruned model and *DarwinLM* on Llama-2-7B and Llama-3.1-8B. *DarwinLM* outperforms uniform models on both pre-trained models, demonstrating the importance of non-uniform sparsity allocation.

Model	Method	Param.	PIQA	WG	ArcC	HS
LLaMA-2-7B	Uniform	3.3B	57.1	53.3	32.2	27.3
	<i>DarwinLM</i>	3.3B	<b>70.0</b>	<b>59.4</b>	<b>36.2</b>	<b>53.5</b>
LLaMA-3.1-8B	Uniform	4.5B	53.6	51.7	23.6	27.1
	<i>DarwinLM</i>	4.6B	<b>69.4</b>	<b>57.3</b>	<b>34.2</b>	<b>44.6</b>
Qwen-2.5-14B-Ins	Uniform	8.6B	72.7	57.6	45.6	27.1
	<i>DarwinLM</i>	8.4B	<b>69.4</b>	<b>57.3</b>	<b>34.2</b>	<b>44.6</b>

Llama-2-7B, the uniformly pruned Llama-3.1-8B model suffers catastrophic degradation. For example, the uniformly pruned model achieves 26.0, 23.6, and 27.1 on ARC-E, ARC-C, and HellaSwag, respectively, close to randomly generated results (25.0%). In contrast, *DarwinLM* significantly improves performance, achieving 59.6, 34.2, and 44.6 on these datasets. Overall, *DarwinLM* shows the best average performance compared to both the uniformly pruned and ZipLM models. After post-compression fine-tuning, *DarwinLM* recovers performance across all benchmarks, with an average score of 63.7. For reference, our model surpasses the open OLMO model with 7B parameters, which is trained with 2.5T tokens. This comparison is meant to signal the fact that, starting from an accurate model, *DarwinLM* can produce competitive models tailored to any runtime/size requirements, at very low training cost.

**Results on Qwen-2.5-14B-Instruct.** To show the scalability of *DarwinLM*, we conduct the pruning on a much larger model, namely Qwen-2.5-14B-Instruct (Qwen, 2024). We prune Qwen-2.5-14B-Instruct with the target level 5. The results are shown in Table 2. First of all, different from Llama-2-7B and Llama-3.1-8B, the uniform pruned model of Qwen-2.5 obtains satisfactory performance on all benchmarks with 56.9 on average, surpassing ZipLM with similar sparsity. This indicates the failure case of ZipLM as

## DarwinLM: Evolutionary Structured Pruning of Large Language Models




Figure 4. Comparison of *DarwinLM* and other one-shot methods that remove modules entirely. Our method consistently outperforms across all sparsity levels, demonstrating the effectiveness of our finer-grained structured pruning approach. Note that the y-axis is log-scaled.

it only optimizes the local error of pruning. However, *DarwinLM* still achieves better than uniform structure. Specifically, *DarwinLM* obtains 59.4 on average on all benchmarks, outperforming the uniform model. After post-compression training with only 2B tokens, the performance of *DarwinLM* increases to 68.1, achieving higher downstream accuracy compared with the reference model, OLMO-0424 trained with 2.05T tokens.

### 4.3. Analysis

**Uniform vs. Non-uniform Pruning.** We now focus on the performance comparison of the uniformly sparse model with the non-uniform model, as shown in Table 3. The one-shot pruning results without post-training are presented for both the uniform model and *DarwinLM* on Llama-2-7B, Llama-3.1-8B and Qwen-2.5-14B-Ins. We apply a target level 5 for all models. Across all benchmarks, the one-shot uniform sparse model performs close to random except Qwen model. In contrast, the non-uniform sparse model, *DarwinLM*, shows significant improvement. For example, on HellaSwag, the uniform model achieves only 27.3 on Llama-2 and 27.1 on Llama-3.1, while *DarwinLM* scores 53.5 and 44.6, respectively. On the ARC-C dataset, *DarwinLM* outperforms the uniform model with scores of 36.2 and 34.2 for Llama-2 and Llama-3.1, compared to 32.2 and 23.6 for the uniform model. These results clearly demonstrate the effectiveness of non-uniform sparsity in structured pruning for large language models and *DarwinLM* can obtain better performance even with higher sparsity compared with uniform sparsity. A full list of results is available in the Appendix.

**Comparison with One-shot Methods under Different Sparsities.** We further compare *DarwinLM* with several current one-shot structured pruning (layer dropping) methods including EvoPress (Sieberling et al., 2024), ShortGPT (Men et al., 2024), Shortened Llama (Kim et al., 2024), and a method based on sliding window cosine similarity (Gromov et al., 2024) on Llama-2-7B. We select 40 samples with 4096 tokens from Fineweb-Edu as the test set and compute the perplexity of each model under different sparsity. For

Table 4. Ablation of our proposed training-aware offspring selection (TAS) on Llama-2-7B with target level 5.

Model	PIQA	SciQ	ArcE
Uniform	57.1	44.1	32.2
<i>DarwinLM</i> w/o TAS	68.8	88.2	63.5
<i>DarwinLM</i>	69.2	88.7	63.8
<i>DarwinLM</i> w/o TAS + 1B token	73.1	91.6	69.0
<i>DarwinLM</i> + 1B token	<b>74.2</b>	<b>92.0</b>	<b>70.8</b>

*DarwinLM*, we choose the target speedups with the closest sparsity settings for comparison. The comparison results are shown in Figure 4. First, we can observe that even if all pruning methods can preserve performance well under the sparsity of 25%, *DarwinLM* still achieves lower perplexity compared with other one-shot pruning methods. Moreover, the performance of ShortGPT and Sliding Window Cosine Similarity shows dramatic degradation after 25% sparsity while the perplexity of Shortened Llama and *DarwinLM* increases slightly before 40%. However, EvoPress also degrades to more than 30 perplexity while *DarwinLM* shows a minor increase for 50% sparsity. Generally, *DarwinLM* outperforms all other one-shot methods under different sparsity settings and keeps stable performance with the increase of sparsity, demonstrating the effectiveness of our method. This is natural also since our method benefits from higher compression granularity.

**Ablation Study.** Finally, we examine the impact of training-aware selection for structure searching and post-training. The results are presented in Table 4. First of all, both models with and without training-aware selection (TAS in the context) searched with 200 generations are better than uniform models. Furthermore, the performance gap of *DarwinLM* with and without TAS is minor before training, indicating that applying TAS generates sparse models with similar performance. However, after 1B token of training for each model, the performance gap between models with and without TAS becomes larger, demonstrating that with training-aware selection, *DarwinLM* is able to select the optimal model for post-training.

## 5. Conclusion

We introduced a novel non-uniform, training-aware structured pruning method called *DarwinLM*, which generates compressed models by evolutionary search. We propose the training-aware offspring selection to select the optimal sparse model for post-training. Experiments on Llama-2-7B, Llama-3.1-8B, and Qwen-2.5-14B-Instruct demonstrate that our approach achieves state-of-the-art performance. *DarwinLM* is remarkably sample-efficient, as it can match or even improve upon the performance of prior methods which required 10x more data and training computation.

## DarwinLM: Evolutionary Structured Pruning of Large Language Models

**Impact Statement**

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

**References**

- Ainslie, J., Lee-Thorp, J., de Jong, M., Zemlyanskiy, Y., Lebrón, F., and Sanghai, S. Gqa: Training generalized multi-query transformer models from multi-head checkpoints. *arXiv preprint arXiv:2305.13245*, 2023.
- An, Y., Zhao, X., Yu, T., Tang, M., and Wang, J. Fluctuation-based adaptive structured pruning for large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024.
- Bisk, Y., Zellers, R., Gao, J., Choi, Y., et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 7432–7439, 2020.
- Clark, C., Lee, K., Chang, M.-W., Kwiatkowski, T., Collins, M., and Toutanova, K. Boolq: Exploring the surprising difficulty of natural yes/no questions. In *NAACL*, 2019.
- Clark, P., Cowhey, I., Etzioni, O., Khot, T., Sabharwal, A., Schoenick, C., and Tafjord, O. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- Dettmers, T., Svirschevski, R., Egiazarian, V., Kuznedelev, D., Frantar, E., Ashkboos, S., Borzunov, A., Hoefler, T., and Alistarh, D. Spqr: A sparse-quantized representation for near-lossless llm weight compression. *arXiv preprint arXiv:2306.03078*, 2023.
- Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Frantar, E. and Alistarh, D. Spdy: Accurate pruning with speedup guarantees. In *International Conference on Machine Learning*, pp. 6726–6743. PMLR, 2022.
- Frantar, E. and Alistarh, D. Sparsegpt: Massive language models can be accurately pruned in one-shot. In *International Conference on Machine Learning*, pp. 10323–10337. PMLR, 2023.
- Frantar, E., Ashkboos, S., Hoefler, T., and Alistarh, D. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*, 2022.
- Gao, L., Tow, J., Abbasi, B., Biderman, S., Black, S., DiPofi, A., Foster, C., Golding, L., Hsu, J., Le Noac’h, A., et al. A framework for few-shot language model evaluation, 12 2023. URL <https://zenodo.org/records/10256836>, 7.
- Groeneveld, D., Beltagy, I., Walsh, P., Bhagia, A., Kinney, R., Tafjord, O., Jha, A. H., Ivison, H., Magnusson, I., Wang, Y., et al. Olmo: Accelerating the science of language models. *arXiv preprint arXiv:2402.00838*, 2024.
- Gromov, A., Tirumala, K., Shapourian, H., Glorioso, P., and Roberts, D. A. The unreasonable ineffectiveness of the deeper layers. *arXiv preprint arXiv:2403.17887*, 2024.
- Gu, Y., Dong, L., Wei, F., and Huang, M. MiniLLM: Knowledge distillation of large language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- Hassibi, B. and Stork, D. Second order derivatives for network pruning: Optimal brain surgeon. *Advances in neural information processing systems*, 5, 1992.
- Hendrycks, D., Burns, C., Basart, S., Zou, A., Mazeika, M., Song, D., and Steinhardt, J. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020.
- Hinton, G., Vinyals, O., and Dean, J. Distilling the knowledge in a neural network, 2015.
- Hsieh, C.-Y., Li, C.-L., Yeh, C.-K., Nakhost, H., Fujii, Y., Ratner, A., Krishna, R., Lee, C.-Y., and Pfister, T. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes. *arXiv preprint arXiv:2305.02301*, 2023.
- Huang, W., Liu, Y., Qin, H., Li, Y., Zhang, S., Liu, X., Magno, M., and Qi, X. Billm: Pushing the limit of post-training quantization for llms. In *Forty-first International Conference on Machine Learning*, 2024.
- Kim, B.-K., Kim, G., Kim, T.-H., Castells, T., Choi, S., Shin, J., and Song, H.-K. Shortened llama: A simple depth pruning for large language models. *arXiv preprint arXiv:2402.02834*, 2024.
- Klein, A., Golebiowski, J., Ma, X., Perrone, V., and Archambeau, C. Structural pruning of large language models via neural architecture search. In *AutoML Conference 2023*, 2023.
- Kurtic, E., Campos, D., Nguyen, T., Frantar, E., Kurtz, M., Fineran, B., Goin, M., and Alistarh, D. The optimal bert surgeon: Scalable and accurate second-order pruning for large language models. *arXiv preprint arXiv:2203.07259*, 2022.

## DarwinLM: Evolutionary Structured Pruning of Large Language Models

- Kurtić, E., Frantar, E., and Alistarh, D. Ziplm: Inference-aware structured pruning of language models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Li, S., Ning, X., Wang, L., Liu, T., Shi, X., Yan, S., Dai, G., Yang, H., and Wang, Y. Evaluating quantized large language models. In *Forty-first International Conference on Machine Learning*, 2024.
- Lin, J., Tang, J., Tang, H., Yang, S., Chen, W.-M., Wang, W.-C., Xiao, G., Dang, X., Gan, C., and Han, S. Awq: Activation-aware weight quantization for on-device llm compression and acceleration. *Proceedings of Machine Learning and Systems*, 6:87–100, 2024.
- Liu, C., Zhao, F., Kuang, K., Kang, Y., Jiang, Z., Sun, C., and Wu, F. Evolving knowledge distillation with large language models and active learning. In Calzolari, N., Kan, M.-Y., Hoste, V., Lenci, A., Sakti, S., and Xue, N. (eds.), *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 2024.
- Liu, J., Cui, L., Liu, H., Huang, D., Wang, Y., and Zhang, Y. Logiqa: A challenge dataset for machine reading comprehension with logical reasoning. *arXiv preprint arXiv:2007.08124*, 2020.
- Lozhkov, A., Ben Allal, L., von Werra, L., and Wolf, T. Fineweb-edu, May 2024. URL <https://huggingface.co/datasets/HuggingFaceFW/fineweb-edu>.
- Ma, S., Wang, H., Ma, L., Wang, L., Wang, W., Huang, S., Dong, L., Wang, R., Xue, J., and Wei, F. The era of 1-bit llms: All large language models are in 1.58 bits. *arXiv preprint arXiv:2402.17764*, 2024.
- Ma, X., Fang, G., and Wang, X. LLM-pruner: On the structural pruning of large language models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- Men, X., Xu, M., Zhang, Q., Wang, B., Lin, H., Lu, Y., Han, X., and Chen, W. Shortgpt: Layers in large language models are more redundant than you expect. *arXiv preprint arXiv:2403.03853*, 2024.
- Muralidharan, S., Sreenivas, S. T., Joshi, R., Chochowski, M., Patwary, M., Shoeybi, M., Catanzaro, B., Kautz, J., and Molchanov, P. Compact language models via pruning and knowledge distillation. *arXiv preprint arXiv:2407.14679*, 2024.
- Qwen, T. Qwen2.5: A party of foundation models, September 2024. URL <https://qwenlm.github.io/blog/qwen2.5/>.
- Sakaguchi, K., Bras, R. L., Bhagavatula, C., and Choi, Y. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.
- Sanh, V. Distilbert, a distilled version of bert: Smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- Sieberling, O., Kuznedelev, D., Kurtic, E., and Alistarh, D. Evopress: Towards optimal dynamic model compression via evolutionary search. *arXiv preprint arXiv:2410.14649*, 2024.
- Tang, S., Ma, L., Li, H., Sun, M., and Shen, Z. Bi-mamba: Towards accurate 1-bit state space models. *arXiv preprint arXiv:2411.11843*, 2024.
- Tao, C., Hou, L., Bai, H., Wei, J., Jiang, X., Liu, Q., Luo, P., and Wong, N. Structured pruning for efficient generative pre-trained language models. In *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 10880–10895, 2023.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Wang, H., Ma, S., Dong, L., Huang, S., Wang, H., Ma, L., Yang, F., Wang, R., Wu, Y., and Wei, F. Bitnet: Scaling 1-bit transformers for large language models. *arXiv preprint arXiv:2310.11453*, 2023.
- Wang, Z., Wohlwend, J., and Lei, T. Structured pruning of large language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 6151–6162, 2020.
- Welbl, J., Liu, N. F., and Gardner, M. Crowdsourcing multiple choice science questions. *arXiv preprint arXiv:1707.06209*, 2017.
- Xia, M., Gao, T., Zeng, Z., and Chen, D. Sheared llama: Accelerating language model pre-training via structured pruning. In *ICLR*, 2024.
- Xiao, G., Lin, J., Seznec, M., Wu, H., Demouth, J., and Han, S. Smoothquant: Accurate and efficient post-training quantization for large language models. In *International Conference on Machine Learning*, pp. 38087–38099. PMLR, 2023.
- Xu, X., Li, M., Tao, C., Shen, T., Cheng, R., Li, J., Xu, C., Tao, D., and Zhou, T. A survey on knowledge distillation of large language models. *arXiv preprint arXiv:2402.13116*, 2024a.

---

**DarwinLM: Evolutionary Structured Pruning of Large Language Models**

---

Xu, Y., Han, X., Yang, Z., Wang, S., Zhu, Q., Liu, Z., Liu, W., and Che, W. Onebit: Towards extremely low-bit large language models. *arXiv preprint arXiv:2402.11295*, 2024b.

Yin, L., Wu, Y., Zhang, Z., Hsieh, C.-Y., Wang, Y., Jia, Y., Pechenizkiy, M., Liang, Y., Wang, Z., and Liu, S. Outlier weighed layerwise sparsity (owl): A missing secret sauce for pruning llms to high sparsity. *arXiv preprint arXiv:2310.05175*, 2023.

Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., and Choi, Y. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.

---

**DarwinLM: Evolutionary Structured Pruning of Large Language Models**


---

## A. Appendix

### A.1. More Implementation Details.

**Details of Second-Order Structured Pruning.** We utilize 2,048 sequences with 4,096 tokens from the Fineweb-Edu dataset as calibration data for Llama-2-7B, Llama-3.1-8B, and Qwen-2.5-14B-Instruct. For pruning efficiency, in the attention module, we fully prune the attention heads, and in the MLP module, we prune entire columns of the output matrix. For Llama-2-7B, we prune the input matrix, as well as the Q, K, and V matrices, based on the pruned output matrix in the attention module. For Llama-3.1-8B and Qwen-2.5-8B-Instruct, we ignore the K and V matrix for pruning. The input and gate matrices in the MLP module are pruned according to the output matrix. On average, pruning both Llama-2-7B, Llama-3.1-8B, and Qwen-2.5-14B-Instruct requires  $4 \times 48$ GB GPU memory. Most of the second-order structured pruning experiments are conducted on a  $4 \times$  NVIDIA L40S machine with 48GB GPU memory.

**Details of Evolutionary Search.** Given a target sparsity level, the searching process starts from the uniform initialization. During selection, we apply 4-steps selection with [1024, 2048, 4096, 8192] token for fitness computation and [10K, 50K, 100K, 200K] token for offspring finetuning. The number of survivors is set as [8, 4, 2, 1] for each step and model. We set the learning rate for offspring training as  $1e-5$ . For llama-2-7B, we apply gradual pruning with target sparsity level 5 at the first stage. We search the structure with 200 steps. After 10B token training, we search again with target sparsity level 6 for 500 steps. For Llama-3.1-8B and Qwen-2.5-14B-Instruct, we search the sparse model with target sparsity level 5 for 200 steps. The search process can be done on a single GPU with 48GB of memory.

Table 5. Hyper-parameter details for post-training on *DarwinLM-2.6B*, *DarwinLM-4.4B*, and *DarwinLM-8.4B*.

Parameter	<i>DarwinLM-2.6B</i>	<i>DarwinLM-4.4B</i>	<i>DarwinLM-8.4B</i>
Learning rate	$1e-4$	$1e-4$	$1e-4$
Global batch size	1,024	1152	2048
Warm-up steps	50 steps	10 steps	50 steps
LR decay scheduler	Cosine	Cosine	Cosine
Context length	4,096	8,192	4,096
Overall tokens	10B	10B	10B

**Details of Post-Training.** We train the final 2.6B sparse model pruned from Llama-2-7B and the 4.4B model pruned from Llama-3.1-8B using 10B tokens for each. Gradient accumulation is applied to achieve a larger global batch size. The models are trained with the Adam optimizer, using a learning rate of  $1e-4$ , and a cosine learning rate decay scheduler to ensure gradual reductions for stable training. No weight decay is applied. The training process is conducted on a cluster of 40 H100 GPUs for 13 hours. Detailed hyperparameters for post-training can be found in Table 5.

### A.2. More Experimental Results.

**More results of the uniform and non-uniform sparse model.** We present a full comparison of uniform and non-uniform sparse models across each benchmark in Table 9. For both Llama-2-7B and Llama-3.1-8B, *DarwinLM* consistently outperforms the uniform pruning approach. Interestingly, we observe that the uniformly pruned Llama-3.1-8B model performs better than *DarwinLM* on the LogiQA dataset. We conjecture this is due to the challenging nature of LogiQA, where even the dense model only achieves a score of 31.1. However, on dense models like Llama-2-7B and Qwen-2.5-14B-Instruct, *DarwinLM* surpasses the uniformly pruned model in performance on LogiQA.

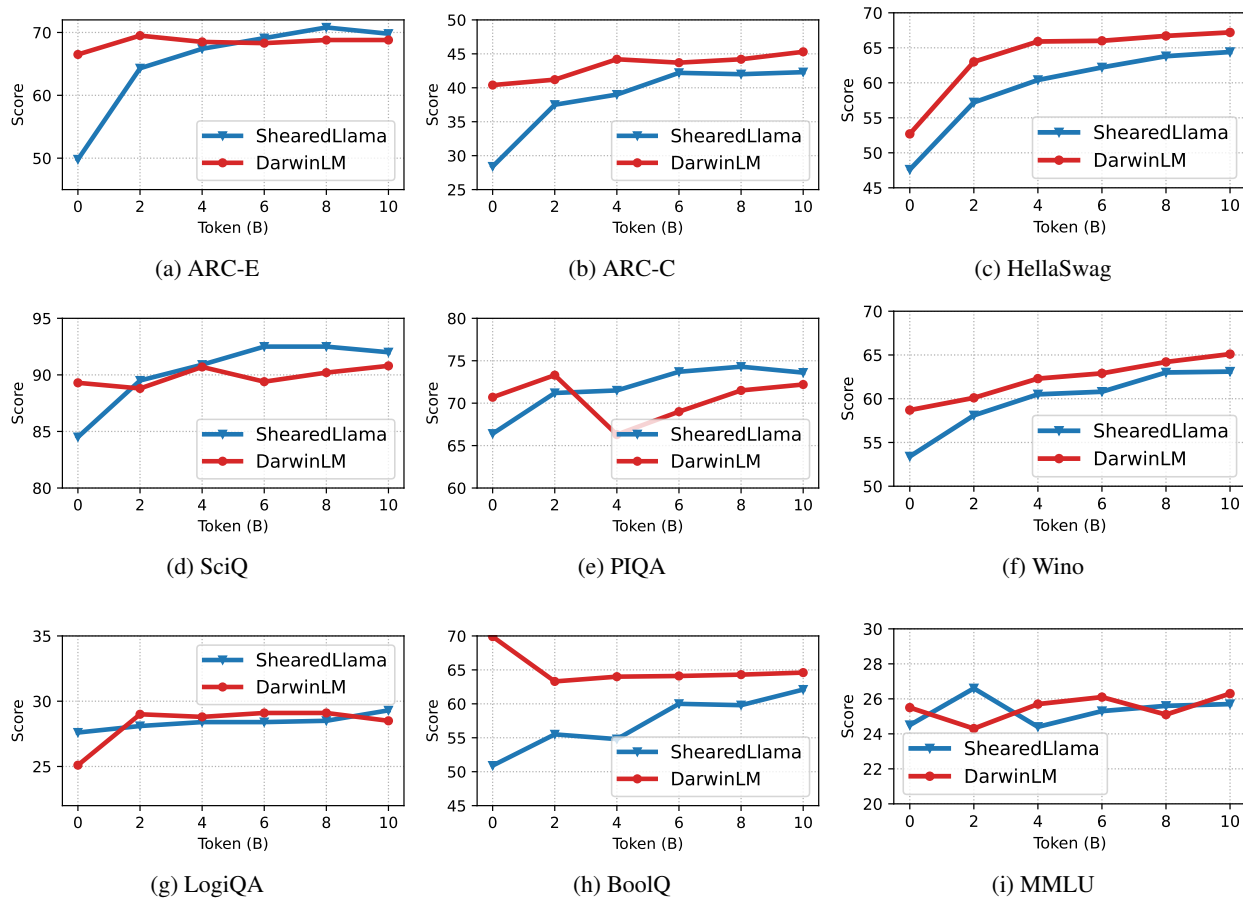
**More Results of Post-Training Comparison with ShearedLlama.** We provide the post-training comparison of ShearedLlama across all benchmarks, with the performance trends for each dataset available in Figure 5. In most cases, *DarwinLM* outperforms ShearedLlama in benchmark evaluations, including SciQ, PIQA, ARC-E, ARC-C, HellaSwag, WinoGrande, LogiQA, BoolQ, and MMLU, demonstrating the effectiveness of our approach.

**Running Time Comparison.** We compare the running time for pruning with ShearedLlama in Table 6. ShearedLlama costs more computation for pruning since it requires additional training to find the mask for the weight. Moreover, the hardware requirement of *DarwinLM* is lower than ShearedLlama, indicating that *DarwinLM* is easy to be adapted to more models.

## DarwinLM: Evolutionary Structured Pruning of Large Language Models

Table 6. Running Time Comparison with ShearedLlama and *DarwinLM*. *DarwinLM* costs less computation compared with ShearedLlama.

Model	Hardware Requirement	Running Time
ShearedLlama	8 × A100-80G	7.4h
DarwinLM	4 × L40s-48G	6.9h

Figure 5. Post-training comparison of ShearedLlama and *DarwinLM* on each benchmark.

## B. Limitations & Future Work

While our non-uniform pruning of LLMs using a speed-aware evolutionary search algorithm offers several advantages, it also has certain limitations: 1) The training-aware algorithm focuses on optimizing both model performance and continued training, but this dual-objective approach may result in suboptimal performance for certain tasks. Focusing on speed can sometimes compromise the accuracy or robustness of the model. 2) The fine-tuning process still requires substantial computational resources, especially when applied to LLMs with billions of parameters. 3) Training-aware selection process introduces extra computation cost, given a large number of offspring. How to find the optimal model for training more efficiently is a direction for future work.

---

**DarwinLM: Evolutionary Structured Pruning of Large Language Models**


---

Table 7. Comparison of *DarwinLM* with uniform pruning on Llama-2-7B.

Method	Param.	SciQ	PIQA	WG	ArcE	ArcC	HS	LogiQA	BoolQ	MMLU	Avg
Dense	6.7B	93.7	78.1	69.3	76.4	53.0	78.6	30.7	82.1	46.6	67.6
Uniform	3.3B	44.1	57.1	53.3	33.5	32.2	27.3	25.0	49.0	23.7	38.4
<i>DarwinLM</i>	3.3B	<b>89.1</b>	<b>70.0</b>	<b>59.4</b>	<b>63.7</b>	<b>36.2</b>	<b>53.5</b>	<b>25.9</b>	<b>65.3</b>	<b>24.8</b>	<b>54.2</b>

Table 8. Comparison of *DarwinLM* with uniform pruning on Llama-3.1-8B.

Method	Param.	SciQ	PIQA	WG	ArcE	ArcC	HS	LogiQA	BoolQ	MMLU	Avg
Dense	8B	96.3	81.2	74.3	81.4	58.2	81.7	31.1	84.0	65.2	72.8
Uniform	4.5B	29.1	53.6	51.7	26.0	23.6	27.1	25.5	62.1	25.7	36.1
<i>DarwinLM</i>	4.6B	<b>84.9</b>	<b>69.4</b>	<b>57.3</b>	<b>59.6</b>	<b>34.2</b>	<b>44.6</b>	24.1	<b>62.2</b>	<b>28.5</b>	<b>51.6</b>

Table 9. Comparison of *DarwinLM* with uniform pruning on Qwen-2.5-14B-Instruct.

Method	Param.	SciQ	PIQA	WG	ArcE	ArcC	HS	LogiQA	BoolQ	MMLU	Avg
Dense	14B	96.8	81.9	79.1	85.7	72.8	85.1	38.5	87.9	80.0	78.6
Uniform	8.6B	78.2	72.7	57.6	76.1	45.6	47.0	28.1	61.6	45.5	56.9
<i>DarwinLM</i> (one-shot)	8.4B	<b>84.3</b>	<b>73.9</b>	<b>60.5</b>	75.7	<b>48.0</b>	<b>53.3</b>	<b>29.3</b>	<b>66.9</b>	<b>43.1</b>	<b>59.4</b>