

# Optimal Transport Depth Up-Scaling

Mingzi Cao   Xi Wang   Nikolaos Aletras

School of Computer Science, University of Sheffield, United Kingdom

{mcao20, xi.wang, n.aletras}@sheffield.ac.uk

## Abstract

Pre-training Large Language Models (LLMs) from scratch at larger scales yields remarkable performance but incurs substantially high training costs. Depth up-scaling provides an efficient alternative by inserting new layers into a pre-trained LLM, avoiding training from scratch. However, most existing methods copying or averaging base layers for new layer, which misalign functionally corresponding neurons, leading to neuron permutation mismatch that harms performance. To address this issue, we propose Optimal Transport Depth Up-Scaling (OT-DUS), which leverages Optimal Transport (OT) theory to align and fuse functionally corresponding neurons module by module in adjacent base layers for new layer construction. OT-DUS achieves better overall performance in both general and specialized domains than existing methods for continual pre-training and supervised fine-tuning across different model sizes and model families. Our analysis of insertion strategies further finds that inserting new layers at higher positions yields not only stronger performance but also improved training efficiency.<sup>1</sup>

## 1 Introduction

Large Language Models (LLMs) performance is largely attributed to scaling laws, where capabilities often improve with increased model and data size (Brown et al., 2020; Kaplan et al., 2020; Wei et al., 2022; Chung et al., 2024). However, scaling poses sustainability challenges, stemming from increased computational (Thompson et al., 2022; Wu et al., 2022; Luccioni et al., 2023) and data (Bender et al., 2021; Birhane et al., 2023; Villalobos et al., 2024) constraints. To address these challenges, model expansion approaches have been proposed, which increase the parameter size of a

<sup>1</sup>Our code is available at <https://github.com/voalmciaf/OPT-DeUS>.




Figure 1: Illustration of our proposed OT-DUS. OT-DUS uses optimal transport to initialize new layers, each derived from two adjacent base layers  $f_i$  and  $f_{i+1}$ . It first aligns each module  $b$  to previous module  $b - 1$  in  $f_i$  (i.e. Alignment within Layer), then aligns it to  $b$  in  $f_{i+1}$  (i.e. Alignment across Layer).

pre-trained model without altering its original architecture. This can be achieved by increasing the number of neurons per layer (i.e., width up-scaling) or the number of layers (i.e., depth up-scaling).

Depth up-scaling has shown better training efficiency and downstream performance than width up-scaling (Du et al., 2024). Recent depth up-scaling approaches update only the new layers for training efficiency while mitigating catastrophic forgetting (Kim et al., 2024; Yang et al., 2025b). Depth up-scaling largely relies on copying (Wu et al., 2024; Kim et al., 2024; Du et al., 2024) or averaging (Yano et al., 2025) for new layer construction. While effective, directly copying or aver-

aging from base layers ignores neuron permutation differences across layers, which can incorrectly align neurons, leading to neuron permutation mismatches that harm performance (Li et al., 2015; Yurochkin et al., 2019a,b). An alternative approach (Yang et al., 2025b) initializes new layers by training an auxiliary neural network, but it introduces additional computational overhead that increases rapidly with model scale. These challenges motivate our main research question: *How can new layers in depth up-scaling be effectively initialized to mitigate neuron permutation mismatches without substantial additional computation?*

Inspired by applying Optimal Transport (OT) (Singh and Jaggi, 2020; Imfeld et al., 2024), we propose Optimal Transport Depth Up-Scaling (OT-DUS). As shown in Figure 1, OT-DUS aligns and fuses adjacent layers module-wise to create neuron-aligned new layers. Newly added layers are initialized via OT and inserted into the top half of the base model. Certain module weights are set to zero for better neuron alignment and function preservation. Our contributions are as follows:

- We introduce OT-DUS, which creates intermediate layer from adjacent layers by neuron alignment via OT. Experiments show that OT-DUS outperforms existing baselines on both continual pre-training and supervised fine-tuning training stages across various model sizes, model families, and diverse tasks.
- OT-DUS achieves the highest overall efficiency, requiring minimal additional computation. Our comprehensive study on layer interpolation position shows that inserting new layers at higher positions leads to higher training efficiency while obtaining better performance.
- Extensive ablation studies validate the contributions of the OT-DUS components, confirming that our design effectively mitigates neuron permutation mismatches. Analysis of neuron functionality reveals a strong correlation between depth up-scaling performance and high transport matrix entropy.

## 2 Related Work

### 2.1 Model Expansion

Model expansion accelerates neural network training by expanding a base model to reduce training time and computational overhead (Chen et al.,

2016; Wei et al., 2016; Chang et al., 2018; Rusu et al., 2022). Network architecture preservation has proven effective for iterative expansion in encoder-only LLMs (Gong et al., 2019; Yang et al., 2020; Chen et al., 2022). More recently, various model expansion approaches have been explored for decoder-only LLMs. Du et al. (2024) showed depth up-scaling yields greater training efficiency and stronger downstream performance compared to width up-scaling. However, prior work primarily focuses on expansion during the pre-training stage with a relatively large pre-training corpus (Shen et al., 2022; Wang et al., 2023, 2024; Yao et al., 2024; Yano et al., 2025), resulting in high overall computational costs. Limited work focuses on post-training expansion (Kim et al., 2024; Wu et al., 2024; Yang et al., 2025b), using a substantially smaller corpus compared to the original pre-training corpus for training efficiency.

### 2.2 Depth Up-Scaling

**Stacking.** Stacking methods insert a series of new layers, typically on top by copying layers from the base model (Du et al., 2024; Kim et al., 2024). Du et al. (2024) proposed stacking entire base layers for stronger downstream performance during pre-training. Kim et al. (2024) introduced SOLAR, a partial stacking approach that omits bottom and top base layers for continual pre-training. However, stacking requires updating the entire model, incurring extra computational costs.

**Interpolation.** Interpolation methods insert new layers inside the base model. Previous work focuses on creating function preservation layers, where the expanded model performs identically to the base model before further training. Achieving function preservation leads to steadier learning processes and better performance. This is achieved by setting the LayerNorm weights to zero for new layer initialization (Shen et al., 2022), initializing the entire new layer to zero (Wang et al., 2024), or employing dynamic masking mechanisms (Yao et al., 2024). Wu et al. (2024) proposed LLaMA PRO, which initializes the inserted new layers by copying weights from the base model. For function preservation, the output matrices of attention and MLP in these new Transformer layers are set to zero, termed zero-init. Yano et al. (2025) initialized new layers by averaging weights from adjacent base layers for pre-training. LESA (Yang et al., 2025b) initializes new layers using an auxiliary network given adjacent layers at interpolation posi-




Figure 2: State-of-the-art depth up-scaling methods and our proposed OT-DUS.

tions as input. However, existing methods largely rely on copying (Kim et al., 2024; Wu et al., 2024) or averaging (Yano et al., 2025) to initialize new layers, neglecting neuron permutation differences.

### 3 Preliminaries

#### 3.1 Depth Up-scaling

Let  $\mathcal{M}$  be a *base* LLM with  $n$  Transformer layers  $\{f_i\}_{i=1}^n$ , parametrized by  $\theta$ . The aim is to obtain an *expanded* model  $\mathcal{M}'$  with parameters  $\theta'$  by introducing  $k$  additional Transformer layers  $\{f'_i\}_{i=1}^k$ .  $\mathcal{M}'$  retains the same layer type and hidden size.

Each Transformer layer  $f_i$  is composed of a sequence of modules. We denote the parameters of layer  $f_i$  by  $\{\mathbf{W}_b^{(i)}\}_{b=1}^B$ , where  $\mathbf{W}_b^{(i)}$  represents the weight matrix of module  $b$  in layer  $f_i$ . Accordingly, each new layer  $f'_i$  is parameterized by  $\{\mathbf{W}_b^{(i)}\}_{b=1}^B$ . **Stacking.**  $\mathcal{M}$  is expanded by adding a set of new layers on top of the base layers to obtain  $\mathcal{M}'$ .  $\circ$  denotes the connection between Transformer layers:

$$\mathcal{M}'(x; \theta') = f'_k \circ \dots \circ f'_1 \circ f_n \circ \dots \circ f_1(x).$$

Each new layer  $f'_i$  is initialized by duplicating a selected base layer  $f_j$  (Du et al., 2024; Kim et al., 2024), where  $j$  is not necessarily equal to  $i$ :

$$\mathbf{W}_b^{(i)} \leftarrow \mathbf{W}_b^{(j)}, \quad \forall b \in \{1, \dots, B\}.$$

**Interpolation.** Figure 2 illustrates different interpolation strategies adopted by existing depth up-

scaling methods.  $\mathcal{M}$  is expanded by inserting new layers between base layers as follows:

$$\mathcal{M}'(x; \theta') = \begin{cases} f'_i \circ f_i, & \text{if a new layer is inserted,} \\ f_i, & \text{otherwise.} \end{cases}$$

Interpolation methods initialize  $f'_i$  from adjacent base layers  $f_i$  and  $f_{i+1}$ . This corresponds to initializing each module weight matrix  $\mathbf{W}_b^{(i)}$  using  $\mathbf{W}_b^{(i)}$  and  $\mathbf{W}_b^{(i+1)}$ . Prior work initialize  $f'_i$  by copying (Wu et al., 2024; Kim et al., 2024; Du et al., 2024), averaging (Yano et al., 2025), or prediction via an auxiliary network (Yang et al., 2025b):

$$\mathbf{W}_b^{(i)} \leftarrow \begin{cases} \mathbf{W}_b^{(i)}, & \text{copying,} \\ \text{Avg}(\mathbf{W}_b^{(i)}, \mathbf{W}_b^{(i+1)}), & \text{averaging,} \\ \text{NN}(\mathbf{W}_b^{(i)}, \mathbf{W}_b^{(i+1)}), & \text{prediction,} \\ \text{OT-DUS}(\mathbf{W}_b^{(i)}, \mathbf{W}_b^{(i+1)}), & \text{our method,} \end{cases}$$

#### 3.2 Optimal Transport

Optimal Transport is a mathematical framework determining the most cost-effective way to transform one probability distribution into another, given a defined transport cost. Formally, let  $\mu = \sum_{i=1}^n \alpha_i \delta(x^{(i)})$  and  $\nu = \sum_{j=1}^m \beta_j \delta(y^{(j)})$  be two discrete probability measures supported on  $\{x^{(i)}\}_{i=1}^n$  and  $\{y^{(j)}\}_{j=1}^m$  with probability distribution  $\alpha = (\alpha_1, \dots, \alpha_n)$  and  $\beta = (\beta_1, \dots, \beta_m)$ , respectively, where  $\delta(s)$  denotes the unit mass at point  $s$ . Given a cost matrix  $\mathbf{C} \in \mathbb{R}^{n \times m}$  where  $\mathbf{C}_{ij}$  is the cost of transporting mass from  $x^{(i)}$  to




Figure 3: Neuron permutation mismatch caused by average and duplicate. Each column represents the order of neurons within layer. Multiple colours of each neuron represent multiple functions it contributes to.

$y^{(j)}$ , the Optimal Transport (OT) problem seeks an transport matrix that minimizes the total cost:

$$\begin{aligned} \text{OT}(\mu, \nu, C) &= \arg \min_{\mathbf{T} \in \mathbb{R}_+^{n \times m}} \sum_{i,j} \mathbf{T}_{ij} C_{ij} \\ \text{s.t. } \mathbf{T} \mathbf{1}_m &= \alpha, \mathbf{T}^\top \mathbf{1}_n = \beta. \end{aligned}$$

## 4 Optimal Transport Depth Up-Scaling

### 4.1 Motivation

Copying and averaging from original layers are commonly used methods for creating new layers (Wu et al., 2024; Kim et al., 2024; Du et al., 2024; Yano et al., 2025). However, such approach neglects neuron permutation mismatch, which is widely present in deep neural networks and Transformers (Li et al., 2015; Yurochkin et al., 2019a,b). During training, a single neuron may contribute to multiple functions (Nguyen et al., 2016), and same-index neurons in different layers may not be functionally corresponding (Sajjad et al., 2022; Klabunde et al., 2025). As shown in Figure 3, averaging from base layers to initialize  $f'_i$  can incorrectly merge neurons with different functionalities, while copying from  $f_i$  to initialize  $f'_i$  breaks the connectivity between  $f_i$  and  $f'_i$ . Thus, directly copying or averaging can cause misalignment between  $f_i$  and  $f_{i+1}$ , potentially harming performance (Li et al., 2015; Yurochkin et al., 2019a,b).

Neuron permutation mismatch can be mitigated by aligning neurons between  $f_i$  and  $f_{i+1}$  using OT, which models functional similarity per neuron. Singh and Jaggi (2020) and Imfeld et al. (2024) show that aligning neurons layer-wise via OT leads to better-initialized new layers for model merging. Recent research further shows that using information from adjacent layers provides stronger initialization than random initialization in depth up-scaling (Du et al., 2024; Yano et al., 2025; Yang et al., 2025b). This inspires our Optimal Transport

Depth Up-scaling (OT-DUS), illustrated in Figure 1. OT-DUS is a progressive interpolation approach that updates only  $f'$ . It aligns and fuses layers  $f_i$  and  $f_{i+1}$  module by module to create  $f'_i$  via OT. OT-DUS inserts new layers  $f'_i$  in the top half of  $\mathcal{M}$ , between base layers  $f_i$  and  $f_{i+1}$ . This layer interpolation strategy provides better performance (Section 7.1) and training efficiency (Section 7.2).

### 4.2 Transport Matrix Flow for OT-DUS

OT-DUS relies on two types of transport matrices:  $\mathbf{T}_{\text{in}}$  and  $\mathbf{T}_{\text{out}}$ . Each module weight matrix  $\mathbf{W}_b^{(i)}$  in  $f'_i$  is assigned a  $\mathbf{T}_{\text{in}}$ .  $\mathbf{T}_{\text{in}}$  aligns  $\mathbf{W}_b^{(i)}$  to  $\mathbf{W}_{b-1}^{(i)}$  within the layer.  $\mathbf{T}_{\text{out}}$  aligns  $\mathbf{W}_b^{(i)}$  to  $\mathbf{W}_b^{(i+1)}$  across layers.  $\mathbf{T}_{\text{in}}$  for  $\mathbf{W}_b^{(i)}$  is initialized by reusing the  $\mathbf{T}_{\text{out}}$  from the previous module  $\mathbf{W}_{b-1}^{(i)}$ .  $\mathbf{T}_{\text{out}}$  is computed by solving an OT problem (Section 4.3).

| Transport Matrix          | Normalization |                                                   | Attention               |                         |                         |                         | MLP                        |                          |                            |
|---------------------------|---------------|---------------------------------------------------|-------------------------|-------------------------|-------------------------|-------------------------|----------------------------|--------------------------|----------------------------|
|                           | Pre-Attn      | Pre-MLP                                           | Query                   | Key                     | Value                   | Output                  | Gate                       | Up                       | Down                       |
| $\mathbf{T}_{\text{in}}$  | $\mathbf{I}$  | $\frac{1}{2}(\mathbf{T}_{\text{O}} + \mathbf{I})$ | $\mathbf{I}$            | $\mathbf{I}$            | $\mathbf{I}$            | $\mathbf{I}$            | $\mathbf{T}_{\text{O}}$    | $\mathbf{T}_{\text{O}}$  | $\mathbf{I}$               |
| $\mathbf{T}_{\text{out}}$ | $\mathbf{I}$  | $\frac{1}{2}(\mathbf{T}_{\text{O}} + \mathbf{I})$ | $\mathbf{T}_{\text{Q}}$ | $\mathbf{T}_{\text{K}}$ | $\mathbf{T}_{\text{V}}$ | $\mathbf{T}_{\text{O}}$ | $\mathbf{T}_{\text{Gate}}$ | $\mathbf{T}_{\text{Up}}$ | $\mathbf{T}_{\text{Down}}$ |

Table 1: Transport Matrix Flow. We manually set  $\mathbf{T}_{\text{in}}$  to each module for alignment within layer.  $\mathbf{T}_{\text{out}}$  is calculated through OT for alignment across layers (except Normalization component). Modules highlighted in underline deviate from Imfeld et al. (2024).

We use Transport Matrix Flow (TMF) to define the assignment of  $\mathbf{T}_{\text{in}}$  for each module (Table 1). Following Imfeld et al. (2024), at the layer entrance of  $f'_i$ ,  $\mathbf{T}_{\text{in}}$  is initialized as the identity matrix  $\mathbf{I}$ . For residual connections (i.e., Pre-MLP Normalization),  $\mathbf{T}_{\text{in}}$  is set by averaging the  $\mathbf{T}_{\text{out}}$  from both paths (i.e. layer entrance and attention output).

However, vanilla TMF is no longer applicable. We therefore modify the TMF design, these modifications are further ablated in Section 7.3.

**Architectural Advancements.** Group-Query Attention (Ainslie et al., 2023) and SwiGLU MLP

(Shazeer, 2020) introduce dimensional mismatches that makes vanilla TMF inapplicable. We thus set  $\mathbf{T}_{\text{in}} = \mathbf{I}$  for Attention Output and MLP Down. As Group-Query Attention and SwiGLU MLP are widely used in current LLMs, including LLaMA (Grattafiori et al., 2024), Qwen (Yang et al., 2025a), and OLMo (OLMo et al., 2025), our modified TMF covers a broad range of mainstream architectures.

**Non-consecutive Alignment.** Alignment is applied only to non-consecutive  $f'_i$ .  $\mathbf{T}_{\text{in}}$  at layer entrance does not carry meaningful cross-layer information. Incorporating this will dilute the impact of alignment. Thus, we set  $\mathbf{T}_{\text{in}} = \mathbf{T}_O$  for MLP Gate and Up that after the residual connection.

### 4.3 Weight Initialization with OT

Given the parameters of layers  $f_i$  and  $f_{i+1}$  and the pre-defined TMF, Algorithm 1 demonstrates how to initialize the new layer  $f'_i$  via OT. The overall procedure consists of five steps, described below.

---

#### Algorithm 1 Optimal Transport Depth Up-Scaling

---

**Input:**  $\mathbf{W}_b^{(i)}, \mathbf{W}_b^{(i+1)}; \mathcal{TMF} = \{\mathbf{T}_{\text{in}}^{(b)}\}_{b=1}^B$  (Table 1)  
**Output:**  $\mathbf{W}'^{(i)}$   
1: **for** base layer  $f_i$  ( $\frac{n}{2} \leq i < n$ ) **do**  
2:   **for** each module  $b$  **do**  
3:     Initialize  $\mu$  and  $\nu$  ▷ OT instantiation  
4:      $\mathbf{C}_{kj} = \|\delta(x^{(k)}) - \delta(y^{(j)})\|_2$  ▷ OT instantiation  
5:      $\mathbf{T}_{\text{in}} \leftarrow \mathcal{TMF}[b]$  ▷ Retrieve  $\mathbf{T}_{\text{in}}$  from  $\mathcal{TMF}$   
6:      $\mathbf{W}_b^{(i)} \leftarrow \mathbf{W}_b^{(i)} \cdot \mathbf{T}_{\text{in}}$  ▷ Alignment within layer  
7:      $\mathbf{T}_{\text{out}} = \text{OT}(\mu, \nu, \mathbf{C})$  ▷ Solve OT problem  
8:      $\mathbf{W}_b^{(i)} \leftarrow \mathbf{T}_{\text{out}}^\top \cdot \mathbf{W}_b^{(i)}$  ▷ Align across layer  
9:      $\mathbf{W}'^{(i)} \leftarrow \frac{1}{2}(\mathbf{W}_b^{(i)} + \mathbf{W}_b^{(i+1)})$  ▷ New Weights  
10:   **end for**  
11:    $\mathbf{W}'_O, \mathbf{W}'_{\text{Down}} \leftarrow \mathbf{0}$  ▷ Zero-init  
12: **end for**

---

**Step-1: OT problem instantiation.** Based on the definition of OT, we initialize  $\mu$  and  $\nu$  over  $\mathbf{W}_b^{(i)}$  and  $\mathbf{W}_b^{(i+1)}$ . Their probability distributions  $\alpha$  and  $\beta$  are initialized uniformly, treating each neurons equally (cf. line 3). To measure neuron differences, we instantiate the unit mass  $\delta$  using weight-based support (Singh and Jaggi, 2020), representing each neuron directly by its weight value without auxiliary constraints. The transport cost  $\mathbf{C}_{kj}$  is defined as the Euclidean distance between the  $k$ -th neuron in  $\mathbf{W}_b^{(i)}$  and the  $j$ -th neuron in  $\mathbf{W}_b^{(i+1)}$  (cf. line 4).  
**Step-2: Alignment within layer.** The permutation change caused by aligning  $\mathbf{W}_{b-1}^{(i)}$  to  $\mathbf{W}_{b-1}^{(i+1)}$  disrupts the original neuron correspondence between  $\mathbf{W}_{b-1}^{(i)}$  and  $\mathbf{W}_b^{(i)}$ . Such permutation change information is stored in  $\mathbf{T}_{\text{in}}$  for  $\mathbf{W}_b^{(i)}$ . To re-

store this,  $\mathbf{W}_b^{(i)}$  needs to align with  $\mathbf{W}_{b-1}^{(i)}$  using  $\mathbf{T}_{\text{in}}$ .  $\mathbf{T}_{\text{in}}$  is defined by  $\mathcal{TMF}$  for each module, shown in Table 1. After retrieving  $\mathbf{T}_{\text{in}}$  (cf. line 5), the alignment within the layer is performed via  $\mathbf{W}_b^{(i)} \leftarrow \mathbf{W}_b^{(i)} \cdot \mathbf{T}_{\text{in}}$  (cf. line 6).

**Step-3: Alignment across layer** We then solve  $\text{OT}(\mu, \nu, \mathbf{C})$  to compute the transport matrix. Imfeld et al. (2024) found that the Sinkhorn-Knopp algorithm (Knight, 2008) is optimal for solving  $\text{OT}(\mu, \nu, \mathbf{C})$  in Transformer fusion. We employ this approach to obtain  $\mathbf{T}_{\text{out}}$  for  $\mathbf{W}_b^{(i)}$  (cf. line 7).  $\mathbf{W}_b^{(i)}$  is then aligned with  $\mathbf{W}_b^{(i+1)}$  using the computed  $\mathbf{T}_{\text{out}}$  via  $\mathbf{W}_b^{(i)} \leftarrow \mathbf{T}_{\text{out}}^\top \cdot \mathbf{W}_b^{(i)}$  (cf. line 8).

**Step-4: Compute new-layer weights.**  $\mathbf{W}'^{(i)}$  is the average of aligned  $\mathbf{W}_b^{(i)}$  and  $\mathbf{W}_b^{(i+1)}$  (cf. line 9).

**Step-5: Zero-Init.** We set  $\mathbf{T}_{\text{in}} = \mathbf{I}$  for  $\mathbf{W}_O$  and  $\mathbf{W}_{\text{down}}$  in TMF due to architectural advancements (Section 4.2), which may cause misalignment problem. Inspired by zero-init (Wu et al., 2024), we set  $\mathbf{W}_O = \mathbf{0}$  and  $\mathbf{W}_{\text{down}} = \mathbf{0}$  (cf. line 11), which naturally resolves this issue while ensuring function preservation, a property crucial for retaining model performance (Wang et al., 2024; Wu et al., 2024).

## 5 Experimental Setup

### 5.1 Base Model

Following prior work (Wu et al., 2024; Kim et al., 2024; Yang et al., 2025b), we use off-the-shelf 32-layer Llama-3.1-8B (Grattafiori et al., 2024) as our *base* model. We further conducted a smaller-scale experiment using off-the-shelf 16-layer Llama-3.2-1B. Following Yang et al. (2025b), we insert a number of new layers equal to 50% of the base layers. The *expanded* model sizes are fixed at 11.5B parameters with 48 layers (adding 16 layers) and 1.72B with 24 layers (adding 8 layers) for all depth up-scaling methods.

### 5.2 Baselines

We experiment with state-of-the-art depth up-scaling methods, as shown in Figure 2. Hyperparameter and implementation details are provided in Appendix A and B.

**Base.** We continue pre-training the *base* model without expansion. All layers are trained.

**SOLAR.** This method copies the bottom and top  $m$  layers from  $\mathcal{M}$  to form  $\mathcal{M}'$ . We choose  $m = 24$  and  $m = 12$  for 11.5B and 1.72B *expanded* models. All layers are trained in line with Kim et al. (2024).

**LLaMA PRO.** It divides  $\mathcal{M}$  into  $g$  groups of  $m$  layers.  $p$  new layers are created by copying the top- $p$  base layers and inserted on top of each group with zero-init. We use  $g = 16$  (11.5B *expanded* models) and  $g = 8$  (1.72B *expanded* models);  $m = 2$  and  $p = 1$  are used throughout. Only  $f'$  are trained following Wu et al. (2024).

**LESA.** This approach uses an auxiliary network to initialize  $f'_i$  given  $f_i$  and  $f_{i+1}$ . LESA inserts  $f'_i$  in the top half of  $\mathcal{M}$ . We insert new layers between  $f_{16}$  and  $f_{32}$  (11.5B *expanded* models), and between  $f_8$  to  $f_{16}$  (1.72B *expanded* models). Only  $f'$  are trained as in Yang et al. (2025b).

### 5.3 Training Data

For Continual Pre-Training (CPT), we opt using data of same size as in Yang et al. (2025b), published after the base model’s knowledge cut-off. We sample 1.5B tokens from the CC-MAIN-2024-51 subset of FineWeb-Edu (Penedo et al., 2024). For supervised fine-tuning (SFT), we choose Alpaca GPT4 (Peng et al., 2023) and update the whole model following Yang et al. (2025b).

### 5.4 Evaluation

Following previous studies (Wu et al., 2024; Yang et al., 2025b), we conduct experiments focusing on general knowledge-related tasks. We further conduct extensive experiments on specialized domains: biomedical (Lee et al., 2019; Gu et al., 2021; Luo et al., 2022; Singhal et al., 2023) and legal (Chalkidis et al., 2019; Zheng et al., 2021; Henderson et al., 2022; Niklaus et al., 2024) as they are widely explored within LLMs. Following common practice in LLM evaluation (Arabelly et al., 2025; Li et al., 2025; Cao et al., 2026), we report results averaged over three random seeds.

**General.** We include ARC-Easy (Clark et al., 2018), LogiQA (Liu et al., 2020), Winogrande (Sakaguchi et al., 2021) for *Reasoning*; CSQA (Talmor et al., 2019), BoolQ (Clark et al., 2019), PIQA (Bisk et al., 2020) for *Commonsense and Knowledge*; MMLU (Hendrycks et al., 2021) for *Examination*; and WikiText (Merity et al., 2017) for *Language Modeling*.

**Biomedical.** Following previous work (Williams et al., 2026), we include MultiMedQA benchmark (Singhal et al., 2022), specifically the PubMedQA (Jin et al., 2019), MedQA (Jin et al., 2021), MedMCQA (Pal et al., 2022) tasks, and relevant subsets from MMLU (Hendrycks et al., 2021) (anatomy,

clinical knowledge, college medicine, medical genetics, professional medicine, college biology).

**Legal.** We follow Williams et al. (2026) in using CaseHOLD (Zheng et al., 2021) and ECtHR (Task A) (Chalkidis et al., 2019) from LexGLUE (Chalkidis et al., 2022) and Legal-MMLU, covering jurisprudence, professional law, and international law (Hendrycks et al., 2021).

## 6 Experimental Results

### 6.1 General Performance

Table 2 (Top) presents the CPT and SFT results of 11.5B *expanded* models. For CPT, we observe that OT-DUS achieves top performance on six out of eight benchmarks, specifically Wiki-PPL (7.73), ARC (82.06), Winogrande (74.77), BoolQ (82.18), PIQA (80.56), MMLU (62.99). This strong performance across various downstream tasks, yielding the highest average score (68.84), highlights the effectiveness. We note that the strong performance continues in SFT. It achieves top performance on Winogrande, CSQA, BoolQ, PIQA, MMLU and second performance on Wiki-PPL, and ARC, with the top average score (69.65).

To analyze performance with different training budgets, we save five checkpoints while training the 11.5B *expanded* models (every 20% training steps), shown in Appendix C. We observe that OT-DUS consistently achieves top performance on at least five out of eight benchmarks across all checkpoints.

### 6.2 Performance at Smaller Scales

Table 2 (Bottom) presents the CPT and SFT results of the 1.72B *expanded* models. For CPT, OT-DUS achieves the highest average score (51.97), ranking first on Wiki-PPL (12.19), LogiQA (22.84), and CSQA (42.97), while ranking second on Winogrande, PIQA, and MMLU. For SFT, OT-DUS again achieves the highest average score, ranking first on CSQA, BoolQ, and PIQA, and second on Wiki-PPL, Winogrande, and MMLU.

Compared with LESA, the strongest baseline on average, OT-DUS achieves top-2 results on more benchmarks in both CPT and SFT settings (Table 2). In CPT, OT-DUS achieves two best and three second-best results at 1.72B, and five best and one second-best results at 11.5B, compared with LESA’s two best and one second-best, and one best and six second-best results, respectively. In SFT, OT-DUS achieves three best and two second-best results at 1.72B, and five best and one second-best

|         |                     | Perplexity ↓                | Zero-shot Performance ↑     |                             |                             |                             |                             |                             |                             |                             |
|---------|---------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|
| Methods |                     | Wiki-PPL                    | ARC                         | LogiQA                      | Wino                        | CSQA                        | BoolQ                       | PIQA                        | MMLU                        | Average                     |
| CPT     | Base-8B             | 8.35 <sub>0.00</sub>        | 79.91 <sub>0.05</sub>       | 26.68 <sub>0.18</sub>       | 72.19 <sub>0.30</sub>       | 65.19 <sub>0.08</sub>       | 81.85 <sub>0.04</sub>       | 79.04 <sub>0.26</sub>       | 58.86 <sub>0.22</sub>       | 66.25 <sub>0.04</sub>       |
|         | SOLAR-11.5B         | 9.90 <sub>0.00</sub>        | 80.02 <sub>0.15</sub>       | 26.88 <sub>0.00</sub>       | 71.80 <sub>0.24</sub>       | 57.36 <sub>0.05</sub>       | 80.76 <sub>0.05</sub>       | 78.51 <sub>0.05</sub>       | 54.32 <sub>0.05</sub>       | 64.24 <sub>0.05</sub>       |
|         | LLaMA PRO-11.5B     | 7.81 <sub>0.00</sub>        | 81.45 <sub>0.15</sub>       | <b>29.39<sub>0.09</sub></b> | 73.93 <sub>0.24</sub>       | 70.71 <sub>0.21</sub>       | 81.71 <sub>0.06</sub>       | 79.89 <sub>0.08</sub>       | 62.59 <sub>0.03</sub>       | 68.52 <sub>0.04</sub>       |
|         | LESA-11.5B          | <u>7.73<sub>0.00</sub></u>  | <u>81.92<sub>0.15</sub></u> | <u>28.06<sub>0.09</sub></u> | <u>74.09<sub>0.12</sub></u> | <b>72.32<sub>0.14</sub></b> | <u>81.89<sub>0.04</sub></u> | <u>80.36<sub>0.06</sub></u> | <u>62.61<sub>0.03</sub></u> | <u>68.75<sub>0.03</sub></u> |
|         | OT-DUS-11.5B (Ours) | <b>7.73<sub>0.00</sub></b>  | <b>82.06<sub>0.06</sub></b> | 27.34 <sub>0.00</sub>       | <b>74.77<sub>0.20</sub></b> | <u>71.96<sub>0.33</sub></u> | <b>82.18<sub>0.07</sub></b> | <b>80.56<sub>0.22</sub></b> | <b>62.99<sub>0.02</sub></b> | <b>68.84<sub>0.04</sub></b> |
| SFT     | Base-8B             | 8.32 <sub>0.00</sub>        | 81.07 <sub>0.09</sub>       | 24.73 <sub>0.15</sub>       | 72.17 <sub>0.28</sub>       | 68.25 <sub>0.25</sub>       | 82.06 <sub>0.08</sub>       | 79.85 <sub>0.12</sub>       | 59.38 <sub>0.19</sub>       | 66.79 <sub>0.05</sub>       |
|         | SOLAR-11.5B         | 9.68 <sub>0.00</sub>        | 80.81 <sub>0.11</sub>       | 25.29 <sub>0.09</sub>       | 71.19 <sub>0.24</sub>       | 60.99 <sub>0.21</sub>       | 81.34 <sub>0.13</sub>       | 79.14 <sub>0.03</sub>       | 55.21 <sub>0.17</sub>       | 64.85 <sub>0.07</sub>       |
|         | LLaMA PRO-11.5B     | 7.81 <sub>0.00</sub>        | 83.30 <sub>0.05</sub>       | <b>27.39<sub>0.18</sub></b> | 73.93 <sub>0.16</sub>       | 72.07 <sub>0.16</sub>       | 82.26 <sub>0.03</sub>       | <u>80.81<sub>0.08</sub></u> | 62.48 <sub>0.14</sub>       | 68.89 <sub>0.03</sub>       |
|         | LESA-11.5B          | <b>7.72<sub>0.00</sub></b>  | <b>83.84<sub>0.04</sub></b> | <u>26.83<sub>0.24</sub></u> | <u>75.61<sub>0.36</sub></u> | <u>72.97<sub>0.29</sub></u> | <u>82.72<sub>0.24</sub></u> | <u>80.52<sub>0.20</sub></u> | <u>63.82<sub>0.22</sub></u> | <u>69.48<sub>0.05</sub></u> |
|         | OT-DUS-11.5B (Ours) | <u>7.73<sub>0.00</sub></u>  | <u>83.80<sub>0.04</sub></u> | 26.73 <sub>0.00</sub>       | <b>75.88<sub>0.44</sub></b> | <b>73.00<sub>0.09</sub></b> | <b>83.31<sub>0.04</sub></b> | <b>80.89<sub>0.11</sub></b> | <b>63.99<sub>0.15</sub></b> | <b>69.65<sub>0.07</sub></b> |
| CPT     | Base-1B             | 13.68 <sub>0.00</sub>       | <u>68.64<sub>0.25</sub></u> | <u>21.56<sub>0.18</sub></u> | 58.43 <sub>0.09</sub>       | 24.57 <sub>0.16</sub>       | 62.69 <sub>0.32</sub>       | 75.03 <sub>0.10</sub>       | 29.27 <sub>0.37</sub>       | 48.60 <sub>0.14</sub>       |
|         | SOLAR-1.72B         | 13.87 <sub>0.00</sub>       | <b>69.00<sub>0.10</sub></b> | 21.35 <sub>0.15</sub>       | 59.54 <sub>0.12</sub>       | 21.54 <sub>0.71</sub>       | 61.12 <sub>0.09</sub>       | 74.88 <sub>0.27</sub>       | 28.70 <sub>0.14</sub>       | 48.02 <sub>0.15</sub>       |
|         | LLaMA PRO-1.72B     | 12.43 <sub>0.00</sub>       | 67.43 <sub>0.17</sub>       | 20.99 <sub>0.09</sub>       | <b>61.48<sub>0.47</sub></b> | 34.23 <sub>0.25</sub>       | <u>62.83<sub>0.10</sub></u> | <b>75.57<sub>0.05</sub></b> | 31.81 <sub>0.10</sub>       | 50.62 <sub>0.09</sub>       |
|         | LESA-1.72B          | <u>12.28<sub>0.00</sub></u> | 66.70 <sub>0.02</sub>       | 21.45 <sub>0.23</sub>       | 59.64 <sub>0.87</sub>       | <u>40.89<sub>0.12</sub></u> | <b>63.72<sub>0.09</sub></b> | 74.74 <sub>0.08</sub>       | <b>33.27<sub>0.21</sub></b> | <u>51.49<sub>0.11</sub></u> |
|         | OT-DUS-1.72B (Ours) | <b>12.19<sub>0.00</sub></b> | 67.10 <sub>0.09</sub>       | <b>22.84<sub>0.23</sub></b> | <u>60.06<sub>0.61</sub></u> | <b>42.97<sub>0.37</sub></b> | 62.75 <sub>0.08</sub>       | <u>75.17<sub>0.14</sub></u> | <u>32.88<sub>0.13</sub></u> | <b>51.97<sub>0.07</sub></b> |
| SFT     | Base-1B             | 13.57 <sub>0.00</sub>       | <u>69.85<sub>0.11</sub></u> | <u>21.97<sub>0.41</sub></u> | 59.35 <sub>0.44</sub>       | 25.83 <sub>0.45</sub>       | 62.94 <sub>0.18</sub>       | 75.72 <sub>0.13</sub>       | 30.17 <sub>0.24</sub>       | 49.40 <sub>0.12</sub>       |
|         | SOLAR-1.72B         | 13.68 <sub>0.00</sub>       | <b>70.41<sub>0.04</sub></b> | <b>22.22<sub>0.09</sub></b> | 59.11 <sub>0.16</sub>       | 25.25 <sub>0.54</sub>       | 60.88 <sub>0.09</sub>       | <u>75.86<sub>0.03</sub></u> | 29.79 <sub>0.35</sub>       | 49.08 <sub>0.13</sub>       |
|         | LLaMA PRO-1.72B     | <b>12.36<sub>0.00</sub></b> | 68.08 <sub>0.06</sub>       | 21.10 <sub>0.23</sub>       | <b>60.62<sub>0.32</sub></b> | 38.19 <sub>0.42</sub>       | 64.09 <sub>0.06</sub>       | <u>76.03<sub>0.21</sub></u> | 30.99 <sub>0.22</sub>       | 51.30 <sub>0.07</sub>       |
|         | LESA-1.72B          | 12.54 <sub>0.00</sub>       | 67.86 <sub>0.11</sub>       | 20.84 <sub>0.09</sub>       | 60.25 <sub>0.60</sub>       | <u>43.73<sub>0.25</sub></u> | <u>65.05<sub>0.20</sub></u> | 75.81 <sub>0.06</sub>       | <b>34.98<sub>0.44</sub></b> | <u>52.65<sub>0.16</sub></u> |
|         | OT-DUS-1.72B (Ours) | <u>12.46<sub>0.00</sub></u> | 68.28 <sub>0.13</sub>       | 21.15 <sub>0.32</sub>       | <u>60.30<sub>0.27</sub></u> | <b>44.69<sub>0.19</sub></b> | <b>65.76<sub>0.07</sub></b> | <b>75.90<sub>0.10</sub></b> | <u>33.22<sub>0.11</sub></u> | <b>52.76<sub>0.04</sub></b> |

Table 2: CPT on 1.5B tokens and SFT (after CPT) performance of 11.5B and 1.72B *expanded* models.

results at 11.5B, compared with LESA’s one best and two second-best, and one best and five second-best results. These results demonstrate the consistent effectiveness of OT-DUS across model scales, training stages, and benchmarks.

### 6.3 Up-scaling Stability at Larger Scales

We follow previous work (Yano et al., 2025; Yang et al., 2025b) by reporting perplexity without any training to evaluate up-scaling stability on larger models. Appendix D presents the perplexity at different model scales. We observe that both LLaMA-Pro and OT-DUS match the base model’s perplexity regardless of model parameters due to zero-init, demonstrating maximum expansion stability compared to other baselines. Surprisingly, we find that LESA’s perplexity sharply increases when applied to Llama-3.2-1B (871.50). We hypothesize this is because smaller models have fewer layers. This leads to less training data for the auxiliary network, consequently causing it to underfit.

### 6.4 Domain Specific Performance

Table 3 presents the performance on biomedical and legal domains. We observe OT-DUS achieves the best overall performance. In biomedical tasks, OT-DUS wins two out of four tasks (i.e. MedMCQA and MedQA), while offering the second best performance on PubMedQA. Strong performance is also observed in the legal domain, where OT-DUS wins two out three legal tasks and achieves second best on the remaining one.

| Methods    |           | MedMCQA                     | MedQA                       | Bio-MMLU                    | PubMedQA                    |
|------------|-----------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|
| Biomedical | Base      | 48.92 <sub>0.28</sub>       | 53.26 <sub>0.00</sub>       | 68.35 <sub>0.10</sub>       | <u>76.33<sub>0.12</sub></u> |
|            | SOLAR     | 45.22 <sub>0.13</sub>       | 47.24 <sub>0.05</sub>       | 57.41 <sub>0.14</sub>       | <u>76.20<sub>0.00</sub></u> |
|            | LLaMA PRO | 54.47 <sub>0.22</sub>       | 59.31 <sub>0.14</sub>       | <u>71.16<sub>0.27</sub></u> | <b>77.67<sub>0.12</sub></b> |
|            | LESA      | <u>56.73<sub>0.19</sub></u> | <u>60.07<sub>0.18</sub></u> | <b>71.40<sub>0.36</sub></b> | 76.00 <sub>0.00</sub>       |
|            | OT-DUS    | <b>56.79<sub>0.16</sub></b> | <b>60.36<sub>0.09</sub></b> | 71.30 <sub>0.29</sub>       | <u>76.33<sub>0.12</sub></u> |
| Methods    |           | Legal-MMLU                  | CaseHOLD                    | ECtHR                       |                             |
| Legal      | Base      | 64.48 <sub>0.01</sub>       | 47.42 <sub>0.06</sub>       | 56.06 <sub>0.12</sub>       |                             |
|            | SOLAR     | 62.62 <sub>0.15</sub>       | 41.02 <sub>0.13</sub>       | 38.90 <sub>0.26</sub>       |                             |
|            | LLaMA PRO | <b>68.58<sub>0.15</sub></b> | 48.31 <sub>0.17</sub>       | 60.40 <sub>0.12</sub>       |                             |
|            | LESA      | 67.14 <sub>0.11</sub>       | <u>51.06<sub>0.06</sub></u> | <u>61.10<sub>0.19</sub></u> |                             |
|            | OT-DUS    | <u>68.14<sub>0.18</sub></u> | <b>51.91<sub>0.09</sub></b> | <b>61.22<sub>0.07</sub></b> |                             |

Table 3: Domain-specific CPT performance of 11.5B *expanded* models.

## 7 Experimental Analysis

### 7.1 Interpolation Positions

We extensively investigate the best interpolation positions of OT-DUS. We evaluate the following strategies: inserting in the bottom half (Btm), in the middle portion (Mid), in the top half (Top), and at the top and bottom quarters (T&B). The layer index ranges are defined as follows:

$$\mathcal{M}'(x; \theta') \circ = \begin{cases} f'_i \circ f_i, & i \leq \frac{n}{2} & \text{if Btm} \\ f'_i \circ f_i, & \frac{n}{4} < i \leq \frac{3n}{4} & \text{if Mid} \\ f'_i \circ f_i, & \frac{n}{2} \leq i < n & \text{if Top} \\ f'_i \circ f_i, & i \leq \frac{n}{4} \text{ or } \frac{3n}{4} \leq i < n & \text{if T\&B} \end{cases}$$

Table 4 illustrates the performance of different interpolation strategies. We observe that OT-DUS-Top is the best performing strategy, overall. It achieves the top performance on five out of eight benchmarks, while obtaining the second best performance on two benchmarks. The performance

difference between interpolation strategies is consistent with previous work, where inserting new layers into the top part offers additional performance gains (Yang et al., 2025b). This further supports previous findings showing that bottom layers in Transformers are more critical (Jawahar et al., 2019), while top layers are less sensitive to modification (Men et al., 2025).

| Interpolation Strategy | Average                     |
|------------------------|-----------------------------|
| OT-DUS-Btm             | 68.39 <sub>0.12</sub>       |
| OT-DUS-Mid             | 68.28 <sub>0.01</sub>       |
| OT-DUS-Top             | <b>68.84<sub>0.04</sub></b> |
| OT-DUS-T&B             | 68.44 <sub>0.04</sub>       |

Table 4: Average performance of OT-DUS-11.5B using different interpolation strategies. Full results are presented in Table 11 (Appendix F).

## 7.2 Training Efficiency

Previous work ignores the impact of interpolation strategy regarding training efficiency (Wu et al., 2024; Yang et al., 2025b). Table 5 shows that progressive depth up-scaling methods considerably outperform SOLAR (22:54:11) in training efficiency. We observe a strong correlation between interpolation positions and efficiency: top-half insertions, exemplified by OT-DUS-Top (12:52:04) and LESA (12:54:07), are notably faster. Conversely, strategies inserting layers in the bottom half, such as OT-DUS-Btm (14:56:00) and LLaMA PRO (14:58:34), require longer training time. This pattern holds regardless of the initialization method. The efficiency differences mainly arise from higher backpropagation costs when updating layers inserted at lower positions in the model.

| Methods          | Trainable | Training Time     |
|------------------|-----------|-------------------|
| SOLAR-11.5B      | 11B       | 22:54:11 (+78.0%) |
| LLaMA PRO-11.5B  | 4B        | 14:58:34 (+16.4%) |
| LESA-11.5B       | 4B        | 12:54:07 (+0.3%)  |
| OT-DUS-Btm-11.5B | 4B        | 14:56:00 (+16.1%) |
| OT-DUS-Mid-11.5B | 4B        | 13:53:14 (+7.9%)  |
| OT-DUS-Top-11.5B | 4B        | 12:52:04          |
| OT-DUS-T&B-11.5B | 4B        | 14:45:38 (+14.7%) |

Table 5: Training time for 11.5B models.

OT-DUS incurs additional computation for solving the OT problem. Compared with baselines that also require extra computation (Table 10, Appendix E), it introduces less overhead and achieves the best overall efficiency among the baselines.

## 7.3 Ablation Studies

| Ablation Variants   |                                                                                   | Average                     |
|---------------------|-----------------------------------------------------------------------------------|-----------------------------|
| TMF                 | Attention Output $\mathbf{T}_{in} = \mathbf{T}_Q$ -1.72B                          | 51.96 <sub>0.03</sub>       |
|                     | MLP Gate and Up $\mathbf{T}_{in} = \frac{1}{2}(\mathbf{T}_O + \mathbf{I})$ -1.72B | 51.79 <sub>0.02</sub>       |
| OT                  | Average-1.72B                                                                     | 51.59 <sub>0.61</sub>       |
|                     | Copy-1.72B                                                                        | 50.13 <sub>0.01</sub>       |
|                     | Random-1.72B                                                                      | 51.39 <sub>0.13</sub>       |
| Zero-Init           | Average + Zero-Init-1.72B                                                         | 51.20 <sub>0.61</sub>       |
|                     | Copy + Zero-Init-1.72B                                                            | 51.67 <sub>0.00</sub>       |
|                     | Random + Zero-Init-1.72B                                                          | 51.58 <sub>0.00</sub>       |
| OT-DUS-1.72B (Ours) |                                                                                   | <b>51.97<sub>0.07</sub></b> |
| TMF                 | Attention Output $\mathbf{T}_{in} = \mathbf{T}_Q$ -11.5B                          | 68.71 <sub>0.02</sub>       |
|                     | MLP Gate and Up $\mathbf{T}_{in} = \frac{1}{2}(\mathbf{T}_O + \mathbf{I})$ -11.5B | 68.73 <sub>0.01</sub>       |
| OT                  | Average-11.5B                                                                     | 68.44 <sub>0.04</sub>       |
|                     | Copy-11.5B                                                                        | 68.59 <sub>0.01</sub>       |
|                     | Random-11.5B                                                                      | 67.11 <sub>0.05</sub>       |
| Zero-Init           | Average + Zero-Init-11.5B                                                         | 68.62 <sub>0.03</sub>       |
|                     | Copy + Zero-Init-11.5B                                                            | 68.82 <sub>0.06</sub>       |
|                     | Random + Zero-Init-11.5B                                                          | 68.43 <sub>0.06</sub>       |
| OT-DUS-11.5B (Ours) |                                                                                   | <b>68.84<sub>0.04</sub></b> |

Table 6: Average performance of ablation variants. Full Results are presented at Table 12 (Appendix G).

**Ablation on TMF design.** We modify the vanilla TMF to account for architectural advancements and non-consecutive alignment (Section 4.2). We further compare our design with alternative variants for architectural advancements (Output:  $\mathbf{T}_{in} = \mathbf{T}_Q$ ) and non-consecutive alignment (Gate/Up:  $\mathbf{T}_{in} = \frac{1}{2}(\mathbf{T}_O + \mathbf{I})$ ). Table 6 shows that OT-DUS achieves the best overall performance (Avg: 51.97), validating our TMF design.

**Ablation on OT alignment.** To validate OT (Section 4.1), we compare it with direct copy, average, and random (Table 6). Compared with the second-best variant, Copy + Zero-Init, OT-DUS outperforms it on 4 of 7 benchmarks at 11.5B and 6 of 7 benchmarks at 1.72B (Table 12), demonstrating consistent gains across model scales and tasks. These results demonstrate that OT alignment provides a more meaningful initialization that mitigates neuron permutation mismatch.

**Ablation on Zero-Init.** OT-DUS adopts zero-init (Section 4.3). We further compare it with non-OT variants using zero-init (Table 6). OT-DUS consistently outperforms these variants, while zero-init improves performance for non-OT variants. Together with the observed up-scaling stability (Section 6.3), these results support the use of zero-init.

**Ablation on Model families.** We further conduct experiments using Qwen-3-8B (Yang et al., 2025a) as *base* model. Table 7 shows that OT-DUS achieves the best overall performance, obtaining the top performance on five out of eight benchmarks. This further validates the generalization of OT-DUS across model families.

| Methods             | Average                     |
|---------------------|-----------------------------|
| Base-8B             | 73.59 <sub>0.01</sub>       |
| SOLAR-11.5B         | 72.36 <sub>0.00</sub>       |
| LLaMA PRO-11.5B     | 73.80 <sub>0.04</sub>       |
| LESA-11.5B          | 73.76 <sub>0.17</sub>       |
| OT-DUS-11.5B (Ours) | <b>74.47<sub>0.06</sub></b> |

Table 7: Average Performance of *expanded* models using Qwen3-8B as *base* model. Full results are presented in Table 13 (Appendix G).

#### 7.4 Analysis of Neuron Functionality

To investigate neuron functionality mapping, we analyze the transport matrix between each new layer and its adjacent base layers. Following Algorithm 1 (Section 4.3), we instantiate the OT problem and solve it with the same hyperparameters, then compute the entropy (Shannon, 1948) of each transport matrix (Table 8). The transport matrix reflects neuron-level functional correspondence, while entropy measures how diffuse the mapping is. Higher entropy indicates a smoother mapping (Cuturi, 2013). SOLAR and LESA are excluded since SOLAR inserts continuous new layers and LESA does not directly use base layers.

| Methods   | Attention |       |       |        | MLP   |       |       |         |
|-----------|-----------|-------|-------|--------|-------|-------|-------|---------|
|           | Query     | Key   | Value | Output | Gate  | Up    | Down  | Average |
| LLaMA PRO | 12.23     | 9.99  | 10.29 | 16.64  | 14.17 | 14.20 | 16.63 | 13.45   |
| Average   | 11.69     | 7.32  | 10.90 | 15.60  | 17.06 | 18.71 | 8.56  | 12.84   |
| OT-DUS    | 13.25     | 10.32 | 11.57 | 16.64  | 16.78 | 18.36 | 16.64 | 14.79   |

Table 8: Layer-averaged Entropy of the transport matrix between base and new layers in the 11.5B *expanded* models. Normalization components are excluded as they do not involve cross-neuron mixing.

We observe a clear correlation between higher transport matrix entropy and stronger performance. OT-DUS achieves the highest entropy (14.79) and the best performance (68.84), followed by LLaMA PRO (13.45/68.52) and Average (12.84/68.44). Higher entropy indicates that information from each neuron is distributed across multiple neurons in the next layer, forming representations through a richer mixture of input features and a more diverse representational space (Tax et al., 2017; Yu et al., 2021), which leads to better performance.

Average yields low entropy for the Key Projection (7.32 vs. 9.99/10.32) and Down Projection (8.56 vs. 16.63/16.64), suggesting these modules are more sensitive to averaging and thus suffer greater performance degradation. Low entropy indicates that functionality from old neurons is concentrated on fewer neurons in the new layer, causing a small subset to dominate the module output. This observation aligns with prior works on the low-rank bottleneck in Query and Key (Bhojanapalli et al., 2020) and parameter redundancy in Down (Pires et al., 2023; Wei et al., 2024).

## 8 Conclusion

We introduced OT-DUS, a depth up-scaling approach that aligns neurons within and across layers via OT to mitigate neuron permutation mismatches. Results demonstrate that OT-DUS achieves better performance in both general and specialized domains with improved efficiency compared to existing approaches across different model sizes, model families, and training stages. Our systematic analysis of interpolation positions reveals that inserting new layers closer to the top leads to higher training efficiency with improved performance. Our neuron functionality analysis further reveals a correlation between strong performance and high transport matrix entropy when initializing new layers.

## Limitations

Following recent work (Kim et al., 2024; Wu et al., 2024; Yang et al., 2025b), we focus on up-scaling already pre-trained LLMs rather than up-scaling during pre-training as in prior work (Shen et al., 2022; Wang et al., 2023, 2024). OT-DUS mitigates neuron permutation mismatch by aligning neurons with similar functionalities across adjacent base layers, which implicitly assumes that base layers have already learned relatively stable functionalities. This assumption may not hold during intermediate stages of pre-training, when neurons are still learning new features. Moreover, applying up-scaling during pre-training would require tremendous training budgets, which are beyond our current computational resources.

## Acknowledgments

We thank Atsuki Yamaguchi for his valuable feedback. We acknowledge IT Services at the University of Sheffield and Bristol Centre for Supercomputing for the provision of HPC services.

## References

- Joshua Ainslie, James Lee-Thorp, Michiel de Jong, Yury Zemlyanskiy, Federico Lebron, and Sumit Sanghai. 2023. [GQA: Training generalized multi-query transformer models from multi-head checkpoints](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4895–4901, Singapore. Association for Computational Linguistics.
- Abhinav Arabelly, Jagrut Nemade, Robert D Nowak, and Jifan Zhang. 2025. [Improving task diversity in label efficient supervised finetuning of LLMs](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 21558–21570, Suzhou, China. Association for Computational Linguistics.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Srinadh Bhojanapalli, Chulhee Yun, Ankit Singh Rawat, Sashank Reddi, and Sanjiv Kumar. 2020. [Low-rank bottleneck in multi-head attention models](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 864–873. PMLR.
- Abeba Birhane, vinay prabhu, Sanghyun Han, Vishnu Boddeti, and Sasha Luccioni. 2023. [Into the LAION’s Den: Investigating Hate in Multimodal Datasets](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 21268–21284. Curran Associates, Inc.
- Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. 2020. [PIQA: Reasoning about Physical Commonsense in Natural Language](#). In *Thirty-Fourth AAAI Conference on Artificial Intelligence*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Jie Cao, Tianwei Lin, Bo Yuan, Rolan Yan, Hongyang He, Wenqiao Zhang, Juncheng Li, Dongping Zhang, Siliang Tang, and Yueting Zhuang. 2026. [MoA: Heterogeneous mixture of adapters for parameter-efficient fine-tuning of large language models](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 21056–21073, San Diego, California, United States. Association for Computational Linguistics.
- Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019. [Neural legal judgment prediction in English](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4317–4323, Florence, Italy. Association for Computational Linguistics.
- Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Katz, and Nikolaos Aletras. 2022. [LexGLUE: A benchmark dataset for legal language understanding in English](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4310–4330, Dublin, Ireland. Association for Computational Linguistics.
- Bo Chang, Lili Meng, Eldad Haber, Frederick Tung, and David Begert. 2018. [Multi-level Residual Networks from Dynamical Systems View](#). In *International Conference on Learning Representations*.
- Cheng Chen, Yichun Yin, Lifeng Shang, Xin Jiang, Yujia Qin, Fengyu Wang, Zhi Wang, Xiao Chen, Zhiyuan Liu, and Qun Liu. 2022. [bert2BERT: Towards Reusable Pretrained Language Models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2134–2148, Dublin, Ireland. Association for Computational Linguistics.
- Tianqi Chen, Ian Goodfellow, and Jonathon Shlens. 2016. [Net2Net: Accelerating Learning via Knowledge Transfer](#). In *International Conference on Learning Representations*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2024. [Scaling Instruction-Finetuned Language models](#). *Journal of Machine Learning Research*, 25(70):1–53.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [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*), pages 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have Solved Question Answering? Try ARC, the AI2 Reasoning Challenge](#). Preprint, arXiv:1803.05457.
- Marco Cuturi. 2013. [Sinkhorn distances: Lightspeed computation of optimal transport](#). In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Tri Dao. 2024. [FlashAttention-2: Faster Attention with Better Parallelism and Work Partitioning](#). In *The Twelfth International Conference on Learning Representations*.
- Wenyu Du, Tongxu Luo, Zihan Qiu, Zeyu Huang, Yikang Shen, Reynold Cheng, Yike Guo, and Jie Fu. 2024. [Stacking Your Transformers: A Closer Look at Model Growth for Efficient LLM Pre-Training](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 10491–10540. Curran Associates, Inc.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2024. [A framework for few-shot language model evaluation](#).
- Linyuan Gong, Di He, Zhuohan Li, Tao Qin, Liwei Wang, and Tieyan Liu. 2019. [Efficient Training of BERT by Progressively Stacking](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2337–2346. PMLR.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and et al. 2024. [The Llama 3 Herd of Models](#). Preprint, arXiv:2407.21783.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021. [Domain-specific language model pretraining for biomedical natural language processing](#). *ACM Trans. Comput. Healthcare*, 3(1).
- Peter Henderson, Mark Krass, Lucia Zheng, Neel Guha, Christopher D Manning, Dan Jurafsky, and Daniel Ho. 2022. [Pile of law: Learning responsible data filtering from the law and a 256gb open-source legal dataset](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 29217–29234. Curran Associates, Inc.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring Massive Multitask Language Understanding](#). In *International Conference on Learning Representations*.
- Moritz Imfeld, Jacopo Galdi, Marco Giordano, Thomas Hofmann, Sotiris Anagnostidis, and Sidak Pal Singh. 2024. [Transformer Fusion with Optimal Transport](#). In *The Twelfth International Conference on Learning Representations*.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. [What Does BERT Learn about the Structure of Language?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. [What disease does this patient have? a large-scale open domain question answering dataset from medical exams](#). *Applied Sciences*, 11(14).
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. [PubMedQA: A dataset for biomedical research question answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2567–2577, Hong Kong, China. Association for Computational Linguistics.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling Laws for Neural Language Models](#). Preprint, arXiv:2001.08361.
- Sanghoon Kim, Dahyun Kim, Chanjun Park, Wonsung Lee, Wonho Song, Yunsu Kim, Hyeonwoo Kim, Yungi Kim, Hyeonju Lee, Jihoo Kim, Changbae Ahn, Seonghoon Yang, Sukyung Lee, Hyunbyung Park, Gyoungjin Gim, Mikyoung Cha, Hwalsuk Lee, and Sunghun Kim. 2024. [SOLAR 10.7B: Scaling Large Language Models with Simple yet Effective Depth Up-Scaling](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 6: Industry Track)*, pages 23–35, Mexico City, Mexico. Association for Computational Linguistics.
- Max Klabunde, Tobias Schumacher, Markus Strohmaier, and Florian Lemmerich. 2025. [Similarity of neural network models: A survey of functional and representational measures](#). *ACM Comput. Surv.*, 57(9).
- Philip A. Knight. 2008. [The Sinkhorn-Knopp Algorithm: Convergence and Applications](#). *SIAM Journal on Matrix Analysis and Applications*, 30(1):261–275.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang.

2019. [Biobert: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*, 36(4):1234–1240.
- Yang Li, Jie Ma, Miguel Ballesteros, Yassine Benajiba, and Graham Horwood. 2025. [Active evaluation acquisition for efficient LLM benchmarking](#). In *Forty-second International Conference on Machine Learning*.
- Yixuan Li, Jason Yosinski, Jeff Clune, Hod Lipson, and John Hopcroft. 2015. [Convergent Learning: Do different neural networks learn the same representations?](#) In *Proceedings of the 1st International Workshop on Feature Extraction: Modern Questions and Challenges at NIPS 2015*, volume 44 of *Proceedings of Machine Learning Research*, pages 196–212, Montreal, Canada. PMLR.
- Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. 2020. [LogiQA: A Challenge Dataset for Machine Reading Comprehension with Logical Reasoning](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3622–3628. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Alexandra Sasha Luccioni, Sylvain Vigui  r, and Anne-Laure Ligozat. 2023. [Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model](#). *Journal of Machine Learning Research*, 24(253):1–15.
- Renqian Luo, Liai Sun, Yingce Xia, Tao Qin, Sheng Zhang, Hoifung Poon, and Tie-Yan Liu. 2022. [Biogpt: generative pre-trained transformer for biomedical text generation and mining](#). *Briefings in Bioinformatics*, 23(6):bbac409.
- Xin Men, Mingyu Xu, Qingyu Zhang, Qianhao Yuan, Bingning Wang, Hongyu Lin, Yaojie Lu, Xianpei Han, and Weipeng Chen. 2025. [ShortGPT: Layers in large language models are more redundant than you expect](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 20192–20204, Vienna, Austria. Association for Computational Linguistics.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer Sentinel Mixture Models](#). In *International Conference on Learning Representations*.
- Anh Nguyen, Jason Yosinski, and Jeff Clune. 2016. [Multifaceted feature visualization: Uncovering the different types of features learned by each neuron in deep neural networks](#). *Preprint*, arXiv:1602.03616.
- Joel Niklaus, Veton Matoshi, Matthias St  rmer, Ilias Chalkidis, and Daniel Ho. 2024. [MultiLegalPile: A 689GB multilingual legal corpus](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15077–15094, Bangkok, Thailand. Association for Computational Linguistics.
- Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafjord, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, Allyson Ettinger, Michal Guerquin, David Heineman, Hamish Ivison, Pang Wei Koh, Jiacheng Liu, Saumya Malik, William Merrill, Lester James V. Miranda, Jacob Morrison, Tyler Murray, Crystal Nam, Jake Ponzanski, Valentina Pyatkin, Aman Rangapur, Michael Schmitz, Sam Skjonsberg, David Wadden, Christopher Wilhelm, Michael Wilson, Luke Zettlemoyer, Ali Farhadi, Noah A. Smith, and Hannaneh Hajishirzi. 2025. [2 olmo 2 furious](#). *Preprint*, arXiv:2501.00656.
- Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2022. [Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering](#). In *Proceedings of the Conference on Health, Inference, and Learning*, volume 174 of *Proceedings of Machine Learning Research*, pages 248–260. PMLR.
- Guilherme Penedo, Hynek Kydl   ek, Loubna Ben al-lal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. 2024. [The Fineweb Datasets: Decanting the Web for the Finest Text Data at Scale](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 30811–30849. Curran Associates, Inc.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. [Instruction tuning with gpt-4](#). *Preprint*, arXiv:2304.03277.
- Telmo Pires, Ant  nio Vilarinho Lopes, Yannick Assogba, and Hendra Setiawan. 2023. [One wide feed-forward is all you need](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 1031–1044, Singapore. Association for Computational Linguistics.
- Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. 2022. [Progressive Neural Networks](#). *Preprint*, arXiv:1606.04671.
- Hassan Sajjad, Nadir Durrani, and Fahim Dalvi. 2022. [Neuron-level interpretation of deep NLP models: A survey](#). *Transactions of the Association for Computational Linguistics*, 10:1285–1303.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. [WinoGrande: an adversarial winograd schema challenge at scale](#). *Commun. ACM*, 64(9):99–106.
- C. E. Shannon. 1948. [A mathematical theory of communication](#). *The Bell System Technical Journal*, 27(3):379–423.
- Noam Shazeer. 2020. [Glu variants improve transformer](#). *Preprint*, arXiv:2002.05202.

- Sheng Shen, Pete Walsh, Kurt Keutzer, Jesse Dodge, Matthew Peters, and Iz Beltagy. 2022. [Staged Training for Transformer Language Models](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 19893–19908. PMLR.
- Sidak Pal Singh and Martin Jaggi. 2020. [Model Fusion via Optimal Transport](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 22045–22055. Curran Associates, Inc.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Nathaneal Scharli, Aakanksha Chowdhery, Philip Mansfield, Blaise Agüera y Arcas, Dale Webster, Greg S. Corrado, Yossi Matias, Katherine Chou, Juraj Gottweis, Nenad Tomasev, Yun Liu, Alvin Rajkumar, Joelle Barral, Christopher Semturs, Alan Karthikesalingam, and Vivek Natarajan. 2022. [Large language models encode clinical knowledge](#). *Preprint*, arXiv:2212.13138.
- Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Le Hou, Kevin Clark, Stephen Pfohl, Heather Cole-Lewis, Darlene Neal, Mike Schaekermann, Amy Wang, Mohamed Amin, Sami Lachgar, Philip Mansfield, Sushant Prakash, Bradley Green, Ewa Dominowska, Blaise Agüera y Arcas, Nenad Tomasev, Yun Liu, Renee Wong, Christopher Semturs, S. Sara Mahdavi, Joelle Barral, Dale Webster, Greg S. Corrado, Yossi Matias, Shekoofeh Azizi, Alan Karthikesalingam, and Vivek Natarajan. 2023. [Towards expert-level medical question answering with large language models](#). *Preprint*, arXiv:2305.09617.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Tycho M.S. Tax, Pedro A.M. Mediano, and Murray Shanahan. 2017. [The partial information decomposition of generative neural network models](#). *Entropy*, 19(9).
- Neil C. Thompson, Kristjan Greenewald, Keeheon Lee, and Gabriel F. Manso. 2022. [The Computational Limits of Deep Learning](#). *Preprint*, arXiv:2007.05558.
- Pablo Villalobos, Anson Ho, Jaime Sevilla, Tamay Besiroglu, Lennart Heim, and Marius Hobbhahn. 2024. [Will we run out of data? Limits of LLM scaling based on human-generated data](#). *Preprint*, arXiv:2211.04325.
- Peihao Wang, Rameswar Panda, Lucas Torroba Henrigen, Philip Greengard, Leonid Karlinsky, Rogério Feris, David Daniel Cox, Zhangyang Wang, and Yoon Kim. 2023. [Learning to Grow Pretrained Models for Efficient Transformer Training](#). In *International Conference on Learning Representations*.
- Yite Wang, Jiahao Su, Hanlin Lu, Cong Xie, Tianyi Liu, Jianbo Yuan, Haibin Lin, Ruoyu Sun, and Hongxia Yang. 2024. [LEMON: Lossless model expansion](#). In *The Twelfth International Conference on Learning Representations*.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. [Emergent Abilities of Large Language Models](#). *Transactions on Machine Learning Research*. Survey Certification.
- Tao Wei, Changhu Wang, Yong Rui, and Chang Wen Chen. 2016. [Network Morphism](#). In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 564–572. PMLR.
- Xiuying Wei, Skander Moalla, Razvan Pascanu, and Caglar Gulcehre. 2024. [Investigating low-rank training in transformer language models: Efficiency and scaling analysis](#). *Preprint*, arXiv:2407.09835.
- Miles Williams, George Chrysostomou, Vitor Amancio Jeronymo, and Nikolaos Aletras. 2026. [Compressing language models for specialized domains](#). In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7393–7415, Rabat, Morocco. Association for Computational Linguistics.
- Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga, Jinshi Huang, Charles Bai, Michael Gschwind, Anurag Gupta, Myle Ott, Anastasia Melnikov, Salvatore Candido, David Brooks, Geeta Chauhan, Benjamin Lee, Hsien-Hsin Lee, Bugra Akyildiz, Maximilian Balandat, Joe Spisak, Ravi Jain, Mike Rabbat, and Kim Hazelwood. 2022. [Sustainable AI: Environmental Implications, Challenges and Opportunities](#). In *Proceedings of Machine Learning and Systems*, volume 4, pages 795–813.
- Chengyue Wu, Yukang Gan, Yixiao Ge, Zeyu Lu, Jiahao Wang, Ye Feng, Ying Shan, and Ping Luo. 2024. [LLaMA Pro: Progressive LLaMA with Block Expansion](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6518–6537, Bangkok, Thailand. Association for Computational Linguistics.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge,

- Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025a. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Cheng Yang, Shengnan Wang, Chao Yang, Yuechuan Li, Ru He, and Jingqiao Zhang. 2020. [Progressively Stacking 2.0: A Multi-stage Layerwise Training Method for BERT Training Speedup](#). *Preprint*, arXiv:2011.13635.
- Yifei Yang, Zouying Cao, Xinbei Ma, Yao Yao, Zhi Chen, Libo Qin, and Hai Zhao. 2025b. [LESA: Learnable LLM layer scaling-up](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 22463–22476, Vienna, Austria. Association for Computational Linguistics.
- Kazuki Yano, Takumi Ito, and Jun Suzuki. 2025. [STEP: Staged Parameter-Efficient Pre-training for Large Language Models](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 374–384, Albuquerque, New Mexico. Association for Computational Linguistics.
- Yiqun Yao, Zheng Zhang, Jing Li, and Yequan Wang. 2024. [Masked Structural Growth for 2x Faster Language Model Pre-training](#). In *The Twelfth International Conference on Learning Representations*.
- Shujian Yu, Kristoffer Wickstrøm, Robert Jenssen, and José C. Príncipe. 2021. [Understanding convolutional neural networks with information theory: An initial exploration](#). *IEEE Transactions on Neural Networks and Learning Systems*, 32(1):435–442.
- Mikhail Yurochkin, Mayank Agarwal, Soumya Ghosh, Kristjan Greenewald, and Nghia Hoang. 2019a. [Statistical Model Aggregation via Parameter Matching](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Mikhail Yurochkin, Mayank Agarwal, Soumya Ghosh, Kristjan Greenewald, Nghia Hoang, and Yasaman Khazaeni. 2019b. [Bayesian Nonparametric Federated Learning of Neural Networks](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 7252–7261. PMLR.
- Lucia Zheng, Neel Guha, Brandon R. Anderson, Peter Henderson, and Daniel E. Ho. 2021. [When does pre-training help? assessing self-supervised learning for law and the casehold dataset of 53,000+ legal holdings](#). In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law, ICAIL '21*, page 159–168, New York, NY, USA. Association for Computing Machinery.

## A Hyper-parameter Details

We set the regularization parameter of Sinkhorn-Knopp algorithm to 0.06, as in Imfeld et al. (2024). We set the global batch size and sequence length to 64 and 2048. For CPT, we use a maximum learning rate of  $1e-4$  for 1.72B *expanded* models and  $5e-5$  for 11.5B *expanded* models. For SFT, the maximum learning rate is set to  $1e-5$  and  $5e-6$ , respectively.

## B Implementation Details

We employ Flash-Attention 2 (Dao, 2024) and mixed-precision bf16 for accelerated training. We use Language Model Evaluation Harness (Gao et al., 2024) for evaluation. 11.5B *expanded* models are trained on four NVIDIA GH200 (96GB) GPUs while 1.72B *expanded* models are trained on a single NVIDIA A100 (80GB). We create all *expanded* models using AMD EPYC 7413 CPU and a single NVIDIA A100 (80GB).

## C Performance across checkpoints




Figure 4: Number of benchmarks that achieve top performance during the training process of 11.5B *expanded* models. Sums may exceed 8 due to ties.

## D Scaling Stability

| Model        | Base  | SOLAR | LLaMA PRO | LESA   | OT-DUS |
|--------------|-------|-------|-----------|--------|--------|
| Llama-3.2-1B | 11.57 | 16.64 | 11.57     | 871.50 | 11.57  |
| Llama-3.1-8B | 7.33  | 9.01  | 7.33      | 9.35   | 7.33   |
| Mistral-24B  | 4.43* | 6.51* | 4.43      | 5.17*  | 4.43   |
| Qwen-2.5-32B | 3.78* | INF*  | 3.78      | 5.67*  | 3.78   |
| Llama-3-70B  | 1.98* | 4.21* | 1.98      | 2.62*  | 1.98   |

Table 9: PPL after 1.5x layer expansion initialization for different *base* models, along with PPL of base models. \* denotes results from Yang et al. (2025b), Results for LLaMA PRO and OT-DUS in bottom-half are obtained via reasonable extrapolation from top-half.

## E Additional Computation Cost

| Expanded Model | Creating Time |
|----------------|---------------|
| LESA-1.72B     | 00:26:15      |
| OT-DUS-1.72B   | 00:02:34      |
| LESA-11.5B     | 04:52:13      |
| OT-DUS-11.5B   | 00:37:16      |

Table 10: Instantiation time for LESA and OT-DUS.

Both LESA and OT-DUS require additional computation. LESA necessitates extracting latent patterns using Singular Value Decomposition (SVD), while OT-DUS requires solving the OT problem module-wise. Table 10 presents the additional computation time required by LESA and OT-DUS to create the *expanded* model. We observe that LESA requires more time compared to OT-DUS (00:26:15 vs. 00:02:34). This time scales massively with larger models (04:52:13 vs. 00:37:16). We hypothesize that this increased time for LESA is mainly caused by the extra computation required for SVD when scaling up base models. Combining Table 5 and Table 10, our OT-DUS achieves the best time efficiency among the baselines.

## F Interpolation Strategy

## G Ablation Results

| Methods    | Perplexity ↓               |                             | Zero-shot Performance ↑     |                             |                             |                             |                             |                             |                             |
|------------|----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|
|            | Wiki-PPL                   | ARC                         | LogiQA                      | Wino                        | CSQA                        | BoolQ                       | PIQA                        | MMLU                        | Average                     |
| OT-DUS-Btm | 7.83 <sub>0.00</sub>       | 81.75 <sub>0.10</sub>       | 28.21 <sub>0.09</sub>       | 74.59 <sub>0.41</sub>       | 70.10 <sub>0.14</sub>       | 81.75 <sub>0.02</sub>       | 79.99 <sub>0.13</sub>       | 62.33 <sub>0.08</sub>       | 68.39 <sub>0.12</sub>       |
| OT-DUS-Mid | <b>7.70<sub>0.00</sub></b> | <b>82.07<sub>0.00</sub></b> | 27.65 <sub>0.00</sub>       | 74.32 <sub>0.05</sub>       | 70.25 <sub>0.23</sub>       | 81.00 <sub>0.12</sub>       | 80.20 <sub>0.09</sub>       | 62.53 <sub>0.06</sub>       | 68.28 <sub>0.01</sub>       |
| OT-DUS-Top | 7.73 <sub>0.00</sub>       | 82.06 <sub>0.06</sub>       | 27.34 <sub>0.00</sub>       | <b>74.77<sub>0.20</sub></b> | <b>71.96<sub>0.33</sub></b> | <b>82.18<sub>0.07</sub></b> | <b>80.56<sub>0.22</sub></b> | <b>62.99<sub>0.02</sub></b> | <b>68.84<sub>0.04</sub></b> |
| OT-DUS-T&B | 7.87 <sub>0.00</sub>       | 81.46 <sub>0.10</sub>       | <b>28.57<sub>0.00</sub></b> | 74.61 <sub>0.18</sub>       | 70.05 <sub>0.05</sub>       | 82.12 <sub>0.02</sub>       | 79.85 <sub>0.03</sub>       | 62.44 <sub>0.03</sub>       | 68.44 <sub>0.04</sub>       |

Table 11: Performance of 11.5B OT-DUS trained on 1.5B tokens using different interpolation strategies.

|           | Methods                                                | Perplexity ↓                |                             | Zero-shot Performance ↑     |                             |                             |                             |                             |                             |                             |
|-----------|--------------------------------------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|
|           |                                                        | Wiki-PPL                    | ARC                         | LogiQA                      | Wino                        | CSQA                        | BoolQ                       | PIQA                        | MMLU                        | Average                     |
| TMF       | Attention Output $T_{in} = T_Q$ -1.72B                 | <b>12.05<sub>0.00</sub></b> | 67.12 <sub>0.02</sub>       | <b>23.30<sub>0.09</sub></b> | 60.25 <sub>0.09</sub>       | 42.64 <sub>0.05</sub>       | 62.89 <sub>0.04</sub>       | 75.39 <sub>0.03</sub>       | 32.13 <sub>0.05</sub>       | 51.96 <sub>0.03</sub>       |
|           | MLP Gate and Up $T_{in} = \frac{1}{2}(T_O + I)$ -1.72B | 12.09 <sub>0.00</sub>       | 66.81 <sub>0.05</sub>       | 22.79 <sub>0.18</sub>       | <b>60.46<sub>0.14</sub></b> | 42.59 <sub>0.00</sub>       | 62.22 <sub>0.02</sub>       | <b>75.48<sub>0.03</sub></b> | 32.20 <sub>0.07</sub>       | 51.79 <sub>0.02</sub>       |
| OT        | Average-1.72B                                          | 12.31 <sub>0.27</sub>       | 67.12 <sub>0.53</sub>       | 22.38 <sub>0.23</sub>       | 59.75 <sub>0.50</sub>       | 41.71 <sub>2.16</sub>       | 62.76 <sub>0.22</sub>       | 74.67 <sub>0.14</sub>       | 32.74 <sub>1.75</sub>       | 51.59 <sub>0.61</sub>       |
|           | Copy-1.72B                                             | 12.62 <sub>0.00</sub>       | <b>68.00<sub>0.02</sub></b> | 22.73 <sub>0.00</sub>       | 59.75 <sub>0.14</sub>       | 34.75 <sub>0.10</sub>       | 62.96 <sub>0.02</sub>       | 74.52 <sub>0.16</sub>       | 28.20 <sub>0.09</sub>       | 50.13 <sub>0.01</sub>       |
|           | Random-1.72B                                           | 16.40 <sub>0.00</sub>       | 66.79 <sub>0.29</sub>       | 22.27 <sub>0.00</sub>       | 59.85 <sub>0.32</sub>       | 40.95 <sub>0.28</sub>       | 62.69 <sub>0.05</sub>       | 74.99 <sub>0.03</sub>       | 32.19 <sub>0.00</sub>       | 51.39 <sub>0.13</sub>       |
| Zero-Init | Average + Zero-Init-1.72B                              | 12.47 <sub>0.27</sub>       | 67.50 <sub>0.54</sub>       | 22.89 <sub>0.40</sub>       | 59.38 <sub>0.72</sub>       | 40.35 <sub>2.12</sub>       | 62.54 <sub>0.32</sub>       | 74.72 <sub>0.08</sub>       | 31.07 <sub>2.30</sub>       | 51.20 <sub>0.61</sub>       |
|           | Copy + Zero-Init-1.72B                                 | 12.23 <sub>0.00</sub>       | 67.09 <sub>0.07</sub>       | 22.27 <sub>0.27</sub>       | 59.48 <sub>0.05</sub>       | 41.01 <sub>0.24</sub>       | <b>64.29<sub>0.02</sub></b> | 75.06 <sub>0.06</sub>       | 32.45 <sub>0.05</sub>       | 51.67 <sub>0.00</sub>       |
|           | Random + Zero-Init-1.72B                               | 12.20 <sub>0.00</sub>       | 66.92 <sub>0.00</sub>       | 21.51 <sub>0.00</sub>       | 59.27 <sub>0.27</sub>       | <b>43.05<sub>0.24</sub></b> | 62.18 <sub>0.02</sub>       | 74.86 <sub>0.09</sub>       | <b>33.29<sub>0.05</sub></b> | 51.58 <sub>0.00</sub>       |
|           | OT-DUS-1.72B (Ours)                                    | 12.19 <sub>0.00</sub>       | 67.10 <sub>0.09</sub>       | 22.84 <sub>0.23</sub>       | 60.06 <sub>0.61</sub>       | 42.97 <sub>0.37</sub>       | 62.75 <sub>0.08</sub>       | 75.17 <sub>0.14</sub>       | 32.88 <sub>0.13</sub>       | <b>51.97<sub>0.07</sub></b> |
| TMF       | Attention Output $T_{in} = T_Q$ -11.5B                 | <b>7.47<sub>0.00</sub></b>  | 81.80 <sub>0.17</sub>       | 27.19 <sub>0.00</sub>       | 73.98 <sub>0.52</sub>       | 71.99 <sub>0.41</sub>       | 82.27 <sub>0.05</sub>       | 80.65 <sub>0.06</sub>       | <b>63.08<sub>0.09</sub></b> | 68.71 <sub>0.02</sub>       |
|           | MLP Gate and Up $T_{in} = \frac{1}{2}(T_O + I)$ -11.5B | <b>7.47<sub>0.00</sub></b>  | 81.80 <sub>0.17</sub>       | 27.19 <sub>0.00</sub>       | 74.06 <sub>0.46</sub>       | <b>72.02<sub>0.41</sub></b> | 82.31 <sub>0.05</sub>       | <b>80.69<sub>0.11</sub></b> | 63.02 <sub>0.08</sub>       | 68.73 <sub>0.01</sub>       |
| OT        | Average-11.5B                                          | 7.95 <sub>0.00</sub>        | 82.07 <sub>0.07</sub>       | 27.60 <sub>0.09</sub>       | 73.69 <sub>0.18</sub>       | 71.20 <sub>0.09</sub>       | 82.11 <sub>0.05</sub>       | 80.23 <sub>0.03</sub>       | 62.16 <sub>0.04</sub>       | 68.44 <sub>0.04</sub>       |
|           | Copy-11.5B                                             | 7.79 <sub>0.00</sub>        | 82.21 <sub>0.02</sub>       | <b>27.65<sub>0.00</sub></b> | 74.24 <sub>0.09</sub>       | 71.94 <sub>0.05</sub>       | 81.60 <sub>0.04</sub>       | 79.74 <sub>0.13</sub>       | 62.78 <sub>0.07</sub>       | 68.59 <sub>0.01</sub>       |
|           | Random-11.5B                                           | 11.69 <sub>0.00</sub>       | 79.71 <sub>0.00</sub>       | 27.19 <sub>0.00</sub>       | 72.51 <sub>0.04</sub>       | 69.45 <sub>0.29</sub>       | 80.57 <sub>0.04</sub>       | 78.58 <sub>0.06</sub>       | 61.78 <sub>0.02</sub>       | 67.11 <sub>0.05</sub>       |
| Zero-Init | Average + Zero-Init-11.5B                              | 7.53 <sub>0.00</sub>        | 82.14 <sub>0.02</sub>       | 27.09 <sub>0.09</sub>       | 74.03 <sub>0.14</sub>       | 71.83 <sub>0.42</sub>       | 82.07 <sub>0.02</sub>       | 80.47 <sub>0.00</sub>       | 62.70 <sub>0.03</sub>       | 68.62 <sub>0.03</sub>       |
|           | Copy + Zero-Init-11.5B                                 | 7.67 <sub>0.00</sub>        | <b>82.36<sub>0.15</sub></b> | 27.65 <sub>0.00</sub>       | <b>74.87<sub>0.05</sub></b> | 71.64 <sub>0.33</sub>       | 82.06 <sub>0.02</sub>       | 80.38 <sub>0.16</sub>       | 62.80 <sub>0.02</sub>       | 68.82 <sub>0.06</sub>       |
|           | Random + Zero-Init-11.5B                               | 7.72 <sub>0.00</sub>        | 81.94 <sub>0.15</sub>       | 26.42 <sub>0.00</sub>       | 73.72 <sub>0.00</sub>       | 71.69 <sub>0.19</sub>       | <b>82.47<sub>0.02</sub></b> | 80.02 <sub>0.06</sub>       | 62.77 <sub>0.01</sub>       | 68.43 <sub>0.06</sub>       |
|           | OT-DUS-11.5B (Ours)                                    | 7.73 <sub>0.00</sub>        | 82.06 <sub>0.06</sub>       | 27.34 <sub>0.00</sub>       | 74.77 <sub>0.20</sub>       | 71.96 <sub>0.33</sub>       | 82.18 <sub>0.07</sub>       | 80.56 <sub>0.22</sub>       | 62.99 <sub>0.02</sub>       | <b>68.84<sub>0.04</sub></b> |

Table 12: Ablation study on 1.72B and 11.5B *expanded* models.

| Methods             | Perplexity ↓               |                             | Zero-shot Performance ↑     |                             |                             |                             |                             |                             |                             |
|---------------------|----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|
|                     | Wiki-PPL                   | ARC                         | LogiQA                      | Wino                        | CSQA                        | BoolQ                       | PIQA                        | MMLU                        | Average                     |
| Base-8B             | <b>9.20<sub>0.00</sub></b> | 84.72 <sub>0.15</sub>       | 39.32 <sub>0.00</sub>       | 73.61 <sub>0.32</sub>       | 82.25 <sub>0.19</sub>       | 82.27 <sub>0.02</sub>       | 79.47 <sub>0.13</sub>       | 73.47 <sub>0.00</sub>       | 73.59 <sub>0.01</sub>       |
| SOLAR-11.5B         | 10.47 <sub>0.00</sub>      | 84.25 <sub>0.12</sub>       | 39.32 <sub>0.00</sub>       | 72.11 <sub>0.18</sub>       | 77.01 <sub>0.05</sub>       | 84.56 <sub>0.05</sub>       | 78.44 <sub>0.06</sub>       | 70.84 <sub>0.03</sub>       | 72.36 <sub>0.00</sub>       |
| LLaMA PRO-11.5B     | 9.24 <sub>0.00</sub>       | <b>85.49<sub>0.05</sub></b> | 37.94 <sub>0.00</sub>       | <b>74.11<sub>0.00</sub></b> | 81.71 <sub>0.19</sub>       | 84.25 <sub>0.00</sub>       | 79.49 <sub>0.09</sub>       | 73.60 <sub>0.03</sub>       | 73.80 <sub>0.04</sub>       |
| LESA-11.5B          | 10.75 <sub>0.00</sub>      | 82.70 <sub>0.07</sub>       | 40.30 <sub>0.09</sub>       | 73.48 <sub>1.09</sub>       | 81.63 <sub>0.05</sub>       | 84.80 <sub>0.05</sub>       | 79.52 <sub>0.03</sub>       | <b>73.88<sub>0.05</sub></b> | 73.76 <sub>0.17</sub>       |
| OT-DUS-11.5B (Ours) | 10.58 <sub>0.00</sub>      | 83.38 <sub>0.00</sub>       | <b>41.83<sub>0.09</sub></b> | 73.48 <sub>0.41</sub>       | <b>84.08<sub>0.09</sub></b> | <b>85.15<sub>0.02</sub></b> | <b>79.54<sub>0.19</sub></b> | 73.84 <sub>0.02</sub>       | <b>74.47<sub>0.06</sub></b> |

Table 13: CPT Performance of *Expanded* models using Qwen3-8B as *base* model.