


Lory: Fully Differentiable Mixture-of-Experts for Autoregressive Language Model Pre-training

Anonymous authors

Paper under double-blind review

Abstract

Mixture-of-experts (MoE) models facilitate efficient scaling; however, training the router network introduces the challenge of optimizing a non-differentiable, discrete objective. Recently, a fully-differentiable MoE architecture SMEAR was proposed (Muqeeth et al., 2023), which ‘softly’ merges experts in the parameter space; nevertheless, its effectiveness was only demonstrated in downstream fine-tuning on classification tasks. In this paper, we present Lory, the first approach that scales such architectures to autoregressive language model pre-training. Lory introduces two key techniques: (1) a causal segment routing strategy that achieves high efficiency for expert merging operations while preserving the autoregressive nature of language models; (2) a similarity-based data batching method that encourages expert specialization by grouping similar documents in training instances. We pre-train a series of Lory models on 150B tokens from scratch, with up to 32 experts and 30B (1.5B active) parameters. Experimental results show significant performance gains over parameter-matched dense models on both perplexity (+13.9%) and a variety of downstream tasks (+1.5%-11.1%). Despite segment-level routing, Lory models achieve competitive performance compared to state-of-the-art MoE models with token-level routing. We further demonstrate that the trained experts in Lory capture domain-level specialization without supervision. Our work highlights the potential of fully-differentiable MoE architectures for language model pre-training and advocates future research in this area.

1 Introduction

Mixture-of-experts (MoE) architectures with sparse activation enable the scaling of model sizes while maintaining high training and inference efficiency (Lepikhin et al., 2020; Fedus et al., 2022; Du et al., 2022; Zoph et al., 2022; Lewis et al., 2021; Zhou et al., 2022). However, training the routing network in MoE architectures introduces the challenge of optimizing a non-differentiable, discrete objective (Shazeer et al., 2017; Zoph et al., 2022). Various techniques—such as switch routing (Fedus et al., 2022), top- k expert-choice routing (Zhou et al., 2022), and linear programming (Lewis et al., 2021)—have been developed to address this challenge, often requiring carefully designed load balancing objectives (Fedus et al., 2022) or introducing additional complexity in assignment algorithms (Lewis et al., 2021; Roller et al., 2021).

Recent research has started to explore fully-differentiable MoE architectures as an alternative to overcome training difficulty. Notably, SMEAR (Muqeeth et al., 2023) is an approach that ‘softly’ merges experts as a weighted average of all the experts’ parameters, as opposed to activating the top- k experts. However, the effectiveness of SMEAR has been only demonstrated in small-scale fine-tuning experiments on downstream classification tasks (Wang et al., 2018). In this work, we propose Lory¹, the first approach that scales such fully-differentiated MoE architectures to autoregressive *language model pre-training*. Unlike text classification tasks which only require routing each input sequence to different experts,

¹Lory is a tribe of parrots with rainbow-like colors, which resembles the spirit of ‘soft’ MoE.

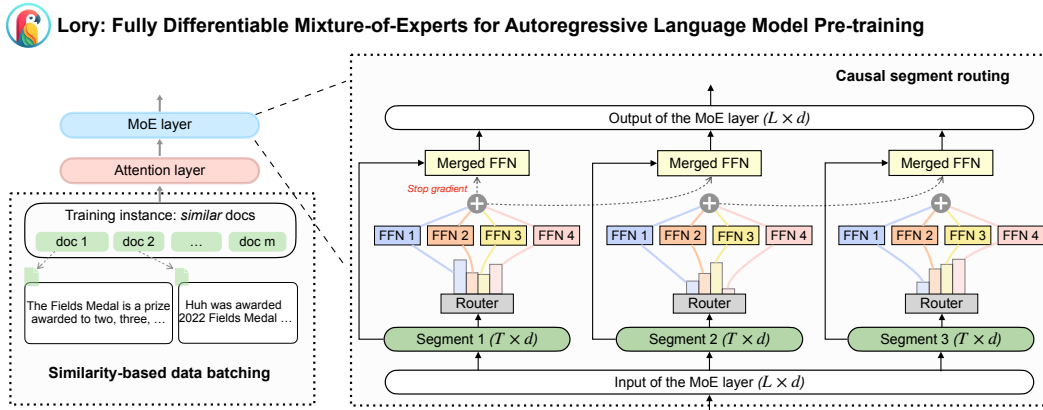


Figure 1: We propose Lory, a fully differentiable MoE architecture designed for auto-regressive language models. Lory is based on two key techniques: First, we propose the *causal segment routing* strategy, which conducts expert merging at the segment level and preserves the autoregressive property of language models. Second, we use the *similarity-based data batching* method to construct training instances, which steers the experts toward specializing in specific domains or topics.

language modeling makes predictions for each input token, and performing token-level routing is prohibitively expensive as the computational cost of merging operations scales linearly with the number of experts.

Lory is based on two key techniques (Figure 1). We first propose *causal segment routing*. For a sequence of input tokens, we split them into multiple segments with a fixed length, and use the previous segment to determine the router’s weights and calculate the merged expert for the subsequent segment. During inference, we can simply use the prompt to make a single routing decision throughout the generation. This segment-level routing strategy preserves the autoregressive nature of language models, while keeping the merging operations efficient. However, since the text data for pre-training language models usually concatenates random sets of documents, we find that such routing can lead to scenarios—*similarity-based data batching* for MoE training, which groups semantically similar documents to form consecutive segments. This idea has been recently proposed to train LMs to better reason across document boundaries (Shi et al., 2023), while we find that it leads to more effective training of expert routing.

We pre-train a series of Lory models from scratch under a training budget of 150B tokens, with 0.3B and 1.5B active parameters, and 8, 16 or 32 experts (up to 6.8B and 29.5B full parameters; see Table 3). Experimental results show that our Lory models significantly outperform equal-sized dense models trained with the same amount of data, achieving performance gains on both perplexity (+13.9%), and a wide range of downstream tasks including commonsense reasoning (+3.7%), reading comprehension (+3.3%), closed-book QA (+1.5%), and text classification (+11.1%). More interestingly, despite that Lory uses segment-level routing, we find it achieves competitive performance compared to state-of-the-art MoE models with *token-level*, non-differentiable discrete routing (Zhou et al., 2022). Our analysis further shows that the trained experts in Lory capture domain-level specialization without any supervision, making it distinct from previous MoE LMs with token-level routing, which only exhibits local patterns uniformly distributed across different domains (Xue et al., 2024; Jiang et al., 2024). Together, we present the first fully-differentiated MoE model that is suitable for language model pre-training, and demonstrate its effectiveness at scale. We hope our work sheds light on the potential of fully differentiable MoE architectures in cultivating specialized experts and we seek to encourage continued exploration in this research field.

2 Preliminaries

2.1 Background: Sparsely-activated MoE

Transformer-based MoE models typically substitute the feed-forward network (FFN) layers with sparsely-activated MoE layers (Shazeer et al., 2017; Fedus et al., 2022; Zoph et al., 2022). Assume that an MoE layer consists of E expert FFNs, each parameterized as $\text{FFN}(\cdot; \theta_1), \dots, \text{FFN}(\cdot; \theta_E)$, where the function $\text{FFN} : \mathbb{R}^{L \times d} \rightarrow \mathbb{R}^{L \times d}$ defines a single expert module. For each token x in a sequence, a MoE layer computes the output y by sparsely activating k experts in this layer and aggregating the outputs through a weighted sum:

$$y = \sum_{i=1}^E e_i \cdot \text{FFN}(x; \theta_i), \quad \text{where } e_i = \text{Top-}k(\text{Softmax}(R(x)))_i. \quad (1)$$

The routing weight e_i for the i -th expert is measured by a *routing network* or *router* R , which takes x as input and calculates the weight for each expert. In practice, to achieve sparsity and computational efficiency, only one (Fedus et al., 2022) or top- k (Lepikhin et al., 2020) experts with the highest routing weights are activated at each MoE layer. The weights of the remaining experts are set to 0 (i.e., $e_i = 0$), eliminating the need to compute $\text{FFN}(x; \theta_i)$ and effectively deactivating the i -th expert.

2.2 Fully Differentiable MoE Architectures via Expert Merging

The primary challenges in training sparsely activated MoE models arise from the difficulties in training discrete routers. A promising direction is to design fully differentiable MoE architectures that do not depend on extra loss formulations for stabilized training. A recent model architecture (Muqeeh et al., 2023) demonstrates the feasibility by calculating a weighted average of all expert FFNs in the parameter space (Matena & Raffel, 2022; Wortsman et al., 2022), thereby creating a "merged FFN". Given an input x and its corresponding routing weights e_i , the output y of a merged FFN is computed as:

$$y = \text{FFN}(x; \sum_i e_i \cdot \theta_i). \quad (2)$$

However, SMEAR has only been evaluated for downstream fine-tuning on text classification tasks, which takes the entire input sequence for routing decisions. Naively extending it to autoregressive language models, which would require computing the merged FFN for each token in a sequence, would be computationally expensive. In this work, we address these challenges by developing a fully differentiable MoE architecture suitable for autoregressive language modeling, and pre-train such models at scale.

3 Our Approach: Lory

In this section, we present Lory, an approach to pre-training fully differentiable MoE language models (Figure 1). The core technique that enables Lory to be fully differentiable is expert merging (Muqeeh et al., 2023, See details in Section 2.2). To make it computationally feasible, we propose a causal segment routing strategy that only merges experts once for each segment, effectively reducing the number of merging operations (Section 3.1). Additionally, we discuss the importance of grouping semantically similar texts to facilitate effective training of the segment-level router (see Section 3.2).

Notations. We denote an input sequence as $X = (x_1, x_2, \dots, x_L)$. By considering a segment size T , we divide the input sequence into $N = \lceil L/T \rceil$ segments, denoted as S_1, S_2, \dots, S_N . We use R to denote the routing network (parameterized as a linear layer) that computes the weights for expert merging. Let h_j represent the hidden representation of the j -th token x_j . The parameters of the i -th expert FFN are denoted by θ_i .

3.1 Efficient Expert Merging via Causal Segment Routing

Challenges. Segment-level routing is an intuitive method to reduce the computational cost of FFN merging by reducing the merging operation from L to N times. However, simply using the current segment to compute the routing weights can cause information leakage.

Training design. We propose *causal segment routing* to effectively route information across segments in an autoregressive manner.² It merges FFNs in an MoE layer based on the previous segment’s information, and uses it to process the current segment. Specifically, given a training instance X that consists of L tokens (e.g., $L = 4096$), we split the training instance into N segments, each of which contains T (e.g., $T = 256$) consecutive tokens. For the k -th segment S_k when $k > 1$, we compute the average of the hidden representations of its preceding segment S_{k-1} , denoted as \bar{h}_{k-1} . Using the average hidden representation allows the model to adapt to prompts of varying lengths during inference. \bar{h}_{k-1} is then utilized to determine the routing weights, resulting in a merged expert $\bar{\theta}$:

$$\bar{h}_{k-1} = \frac{1}{T} \sum_{x \in S_{k-1}} h_x, \quad e_i = \text{Softmax}(R(\bar{h}_{k-1})), \quad \bar{\theta} = \sum_i e_i \cdot \theta_i. \quad (3)$$

We then use the merged expert $\bar{\theta}$ to process all the tokens in the current segment S_k (See Equation 2). This approach guarantees that the routing decisions made by the model are based exclusively on data from preceding positions. For the first segment S_1 , the representation of the segment itself is used to compute the merging weights for its own FFN. To prevent information leakage, we implement a stop-gradient operation on $R(H_1)$. As demonstrated in Appendix B, merging experts at the segment level incurs minimal overhead compared to the training of dense models.

Prompt-only routing during inference. During inference, we begin with a given prompt and make a single routing decision based on the average hidden representations of the prompt. This routing decision determines a merged FFN and it is used consistently throughout the entire generation process. It is important to note that this inference process is as simple and computationally efficient as dense models.³

3.2 Similarity-based Data Batching

Ensuring high similarity for segments in a single training instance is crucial for effectively training Lory. The common approach of randomly concatenating documents to form training instances leads to less specialized experts, impairing model performance (see Figure 4). To mitigate this issue, we employ a similarity-based data batching technique inspired by Shi et al. (2023). We measure document similarity using Contriever (Izacard et al., 2021) and sequentially concatenate similar documents to construct training instances based on a greedy search algorithm (see Appendix C). Although we employ a data batching technique similar to Shi et al. (2023), it is worth noting the difference in motivation. While their work aimed to improve language models’ reasoning across document boundaries, we find this technique effective in encouraging expert specialization within MoE models.

4 Experiments

In this section, we evaluate Lory by training a series of language models from scratch. We first describe the experimental setups (Section 4.1) and then present the results (Section 4.2).

²A piece of pseudocode of the *causal segment routing* strategy can be found in Appendix A.

³In Appendix G.3, we compare the prompt-only routing strategy to using the causal segment routing strategy that faithfully follows the training design. We find that these two strategies do not lead to significant differences in performance on downstream tasks. Given the efficiency advantage of making a single routing decision, we adopt the prompt-only strategy as the default approach.

4.1 Setups

Models. We evaluate our approach by training decoder-only Transformer models which consists of active parameters of 0.3B and 1.5B.⁴ For each FFN layer in the Transformer model, we replace it with MoE layers with E ($E \in \{8, 16, 32\}$) experts with exactly the same architecture. Appendix D shows the configuration of model architectures as well as the total parameter count. We follow LLaMA (Touvron et al., 2023a) and use SwiGLU (Shazeer, 2020) as the activation function in FFNs. We use the same tokenizer as the LLaMA models (Touvron et al., 2023a;b). All models are trained with a 4096-token context window. In the causal segment routing strategy, we set the length of each segment to be $T = 256$.

Training details. We employ the AdamW optimizer (Loshchilov & Hutter, 2017) with $\beta_1 = 0.9$ and $\beta_2 = 0.95$ and use a learning rate of $2e - 4$ with a cosine learning rate scheduler. All models with a batch size of 1 million tokens. We employ the data parallelism with the ZeRO optimization (Rajbhandari et al., 2020) for the distributed training.⁵ At the beginning of training, we train a parameter-matched dense model and duplicate the FFN layers as initialization of the MoE model. In our experiments, we use the first 5% training steps as the warmup to initialize the MoE weights. We find that without warmup training, there may be more experts under-utilized (see Appendix G.4 for an ablation study). We also apply a linear warmup to the learning rate scheduler for the first 5% training steps.

Training datasets. We randomly sample a subset of the Commoncrawl dataset (Wenzek et al., 2019) as the training data. The full training dataset consists of 150 billion tokens in total. We apply the similarity-based data batching method on this subset of construct all the training instances, following Shi et al. (2023). See Appendix C for details of the data batching method.

Evaluation datasets. We evaluate all the models on language modeling tasks by measuring the perplexity of trained models on held-out evaluation datasets sampled from arXiv, Books Corpora, Wikipedia, C4 (Raffel et al., 2020), and Python code (a Python subset of Github). Each evaluation dataset contains 1K samples, each of which consists of 4096 tokens.

We also evaluate models in downstream tasks with in-context learning (Brown et al., 2020), including common sense reasoning: BoolQ (Clark et al., 2019), PIQA (Bisk et al., 2020), SIQA (Sap et al., 2019), HellaSwag (Zellers et al., 2019), WinoGrand (Sakaguchi et al., 2021); reading comprehension: RACE (Lai et al., 2017), ARC (Clark et al., 2018); closed-book QA: Natural Questions (Kwiatkowski et al., 2019), TriviaQA (Joshi et al., 2017); and text classification: AGNews (Zhang et al., 2015), SST-2 Socher et al. (2013), Amazon and Yelp (Zhang et al., 2015), FEVER (Thorne et al., 2018), MRPC (Dolan & Brockett, 2005). For text classification tasks, we follow the evaluation setup of Min et al. (2022); for the rest of tasks, we follow the same setup as Touvron et al. (2023b).

4.2 Main Results

Training efficiency and convergence. Figure 2 (left) shows the training loss curves of the dense model and our MoE models with different model sizes. First, we find that with the same amount of training tokens, our models clearly achieve better training loss compared to the dense model baseline. For the 0.3B and 1.5B models, our models with 32 experts achieve the same level of loss with fewer than half of the training tokens. This indicates that our approach achieves much better performance with the same training compute (see analysis of additional FLOPs from MoE layers in Appendix B). We also observe that when using more experts, we are able to gain more improvement.

⁴In Appendix E, we additionally conduct experiments on a 7B dense model and a 7B/4E MoE model *without* using similarity-based data batching. Due to the limited computing resources, we are not able to train 7B models on the batched dataset.

⁵In Section H.1, we discuss more training strategies when scaling up model sizes (e.g., > 100B).




Figure 2: Left: training curves (log perplexity) of models with different sizes and experts. Right: Perplexity of trained models on different evaluation sets (arXiv, Books, Wikipedia, C4, and Python). We include the detailed model configurations and sizes in Appendix D.

Model	Commonsense Reasoning					Reading Comprehension			
	PIQA	SIQA	BoolQ	HellaSwag	WinoGrande	RACE-m	RACE-h	ARC-e	ARC-c
0.3B	65.8	42.7	44.6	34.6	51.2	41.7	30.9	51.5	21.3
0.3B/8E	67.5	41.2	41.2	34.8	54.4	43.1	31.4	52.4	22.1
0.3B/16E	67.2	44.1	56.6	34.9	54.1	43.9	31.1	54.8	24.9
0.3B/32E	68.2	43.0	58.0	34.7	53.4	42.7	32.0	57.4	26.3
1.5B	71.2	45.0	54.0	43.9	60.9	50.1	36.7	65.0	31.0
1.5B/8E	72.1	45.2	62.0	43.6	63.7	51.2	36.5	66.3	32.5
1.5B/16E	71.3	45.0	56.0	43.7	61.5	51.7	37.3	66.3	32.7
1.5B/32E	72.1	47.1	59.9	43.8	61.9	51.5	32.4	66.7	32.7

Model	Closed-book QA		Text Classification						Avg
	NQ	TQA	AGNews	Amazon	SST-2	Yelp	Fever	MRPC	
0.3B	4.7	8.8	30.3	53.6	54.6	66.0	47.6	62.0	41.8
0.3B/8E	5.3	9.0	38.4	52.3	54.6	62.6	56.6	59.0	42.7
0.3B/16E	6.0	10.2	36.3	75.6	53.3	64.0	57.0	65.0	45.8
0.3B/32E	5.3	10.2	47.3	64.0	55.3	73.3	55.7	56.0	46.0
1.5B	7.6	23.8	64.0	65.3	80.0	58.6	59.0	66.7	51.9
1.5B/8E	7.3	24.2	65.0	94.0	80.0	88.3	57.0	64.0	56.1
1.5B/16E	7.3	25.6	61.6	78.3	84.6	93.6	57.3	63.6	55.1
1.5B/32E	7.0	25.4	62.3	94.7	85.0	95.3	56.3	66.7	56.5

Table 1: We compare the Lory MoE models with the parameter-matched dense models on downstream tasks, including commonsense reasoning, reading comprehension, closed-book QA, and text classification.

Language modeling. We evaluate trained models on language modeling evaluation sets. As shown in Figure 2 (right), our MoE models outperform the dense baseline in all domains, significantly reducing perplexity. For example, our 0.3B/32E model achieves a relative improvement of 13.9% on Books compared to the 0.3B dense model. We observe that the improvement is especially large in test domains that are markedly different from the domains of the training dataset (e.g. Python). We consider this as a strong indication of expert specialization in specific domains (We further study expert specialization in Section 5.4).

Downstream tasks. Table 1 shows the model performance on downstream tasks. We observe significant performance across all tasks. For example, our 0.3B/32E model achieves an average performance improvement of +3.7% in common sense reasoning, +3.3% in reading comprehension, +1.5% in reading comprehension, and +11.1% in text classification.

5 Analysis and Ablation Studies

In this section, we conduct ablation studies and analysis to understand the essence of each component of our approach.

5.1 Importance of Causal Segment Routing

We compare our causal segment routing strategy with an alternative *prefix routing* strategy for training. In prefix routing, expert merging is performed only once for each sequence based on the first segment. The merged FFN is then used to process the rest of the sequence without further updates. Figure 3 shows that using only a prefix for routing leads to much worse performance compared to causal segment routing. These results highlight the importance of using every segment to provide strong training signals for routers.




Figure 3: Training curves of causal segment routing and prefix routing. The latter is a straightforward segment-level routing strategy that uses the first segment to route the entire input.




Figure 4: Left: Training curves of similarity-based data batching (*sim batch*) or the standard random batching (*rand batch*). Right: Training loss difference between Lory and a dense model when using different batching strategies. Lory leads to a larger loss improvement over the dense model when using similarity-based data batching.

5.2 Importance of Similarity-based Data Batching

To investigate the importance of similarity-based data batching, we compare the performance improvement of MoE models over dense models with and without this batching method. Figure 4 (left) shows the training loss of dense (0.3B) and MoE models with eight experts (0.3B/8E) using similarity-batched (*sim batch*) and randomly-batched (*rand batch*) data. MoE models consistently outperform dense models in both setups. However, the loss improvement (i.e., the difference in loss between dense and MoE models) is much larger with similarity-based batching, and this effect is amplified with more training data (Figure 4, right). These results strongly support the importance of similarity-based batching for effectively training our MoE model.

5.3 Comparison with Existing MoE Models

We compare our approach with Expert Choice (EC) (Zhou et al., 2022), a state-of-the-art MoE method that ensures balanced load during training by having each expert select top- k inputs according to the routing weights. We consider two variants of EC MoE models, both with a capacity factor of 1 to match the computation of our MoE models. First, we train a sparse EC MoE model using our segment routing strategy, where each expert selects top segments and processes all tokens within those segments. This variant allows us to directly compare our expert-merging strategy with the expert choice method while using the same segment-level routing approach. Second, we consider the original EC setting with token-level routing to provide an end-to-end comparison with state-of-the-art MoE models using the same amount of training computation. Figure 5 shows the training loss curves. We observe that Lory (blue curve) significantly outperforms segment-level




Figure 5: Comparison with the state-of-the-art MoE training technique Expert Choice (EC) with a segment-level or token-level routing. For both EC models, we use the capacity factor of 1 with the same amount of FLOPs as our training method for the fair comparison.

EC (orange curve) with the same routing setting, suggesting that a fully differentiable architecture is more effective than a sparse MoE when using the same routing strategy. Comparing Lory with the token-level EC model (green curve), we find that Lory achieves competitive results despite using segment-level routing and not requiring any advanced training techniques. These results highlight the significant potential of Lory. In Appendix G.1, we compare Lory and EC on held-out evaluation sets. We find Lory achieves much better perplexity compared to the token-level EC model, while performing similarly on other domains (arXiv, Books, Wiki, C4). Our analysis in Section 5.4 demonstrates that Lory learns experts specialized in specific domains (e.g., Python code), potentially improving performance in less frequent domains.

5.4 Expert Utilization and Specialization

Utilization: How many experts are actively utilized? One potential issue of training MoE models is the models may collapse to dense models because most experts are under-utilized (e.g., some experts have never been activated). In Appendix G.2, we show although without using any auxiliary loss on load balancing, Lory is able to achieve high expert utilization, preventing the MoE models from collapsing to dense models.

Specialization: What do experts learn? In order to study the expert specialization, we investigate the averaged routing weights at different layers of the 0.3B/8E model, on different domains (Books, arXiv, Python, and Wikipedia). Figure 6 shows the routing weights at layer 0, 11, and 23 (the first, middle, and last layer) of the 0.3B/8E model.⁶ First, we find that there exists clear domain-level expert specialization in our trained MoE models, even though no additional domain-level supervision is used during training. For instance, expert 7 at layer 11 is specialized to process inputs in the arXiv domain. We also observe that routing weights on arXiv and Python code are more similar compared to Books and Wikipedia, likely because LaTeX code and Python code are dissimilar to natural language. Second, experts at the middle or high layers are more specialized in specific domains, while the routing weights at lower layers are similar and flat across domains.




Figure 6: Averaged routing weights at layer {0, 11, 23} of the 0.3B/8E model on different domains (Books, arXiv, Python, Wikipedia). We observe that the experts in our MoE models learn domain-level specialization, especially at middle and higher layers.

It is worth noting that our learned experts behave differently from those of prior token-level MoE models, where shallow token-level specialization is observed. For example, some experts are specialized for a specific type of word (e.g., punctuations, articles), and few deep semantic features are captured by the learned routers (Jiang et al., 2024; Lewis et al., 2021; Zoph et al., 2022; Shazeer et al., 2017; Xue et al., 2024). Our models learn domain-level specialization, which we attribute to the segment-level routing strategy used during training. This strategy allows routers to capture global semantic features beyond the token level. The complementary nature of features captured by segment/sentence-level and token-level routing strategies suggests the possibility of combining them to build even stronger models, and we leave it for future work.

⁶In Appendix F, we show the averaged routing weights at all layers of the 0.3B/8E model.

5.5 More Analysis

In Appendix G, we further show that (1) during inference of downstream tasks, routing the entire input prompt once or routing each segment does not make substantial differences on the tasks we evaluate; (2) warmup training is crucial to achieve high expert utilization, especially when training MoE models with a large number of experts. In addition, we discuss the possibility of converting Lory to sparse models for more efficient inference in Appendix H.2.

6 Related Work

Mixture of Experts. Sparsely activated MoE models (Shazeer et al., 2017) have been proposed to demonstrate the potential of massively scaling up model sizes. GShard (Lepikhin et al., 2020) adapts the sparse MoE architecture into Transformer models and achieves strong results on machine translation. Recent work has extended it to general language models (Fedus et al., 2022; Zoph et al., 2022; Jiang et al., 2024; Dai et al., 2024; Zhou et al., 2022; Du et al., 2022; Artetxe et al., 2022; Xue et al., 2024). Traditional MoE models are trained to route given inputs to one or a few specialized expert modules, which introduces a non-differentiable, discrete decision-learning problem. These existing models are trained with the top-1 or top-2 routing strategy on a carefully designed load balancing objective (Lepikhin et al., 2020; Fedus et al., 2022; Zoph et al., 2022), or employ complicated assignment algorithms to distribute inputs (Lewis et al., 2021; Roller et al., 2021; Zhou et al., 2022). Training MoE models has been shown to be difficult, facing the issues of training instability, expert under-specialization, poor training efficiency (Zoph et al., 2022).

Our approach enables end-to-end gradient back-propagation by employing fully differentiable MoE architectures. SMEAR (Muqeeth et al., 2023) proposes softly merging experts by taking a weighted average on the parameter space. However, SMEAR is only applied to text classification tasks with an encoder backbone. Although Lory shares a similar expert merging technique, it is the first approach that scales such architecture to autoregressive language model pre-training. Soft MoE (Puigcerver et al., 2023) is another fully-differentiable MoE architecture which enables end-to-end gradient back-propagation. However, it is only evaluated on vision tasks and does not apply to autoregressive language model pre-training either. We leave how to extending Soft MoE on decoder language models as the future work.

Similarity-based data batching. There exists research that applies a similar data batching method during training. In-context pre-training (Shi et al., 2023) groups relevant documents together to encourage language models to leverage long-range contexts and improve the results of in-context learning and retrieval augmentation. Zhong et al. (2022) batch documents with high lexical similarity to collect more positive pairs in a contrastive learning framework to provide stronger training signals. Although sharing the same idea, the goal of our data batching method is to avoid routing irrelevant documents together, which may hurt the expert specialization.

7 Conclusion

In this paper, we propose Lory, a fully-differentiable MoE designed for autoregressive language model pre-training. Our extensive experiments demonstrate that Lory significantly outperforms its dense counterpart on language modeling tasks and downstream applications. We also observe that trained experts are highly specialized and capable of capturing domain-level information. Future research includes further scaling up Lory, combining token-level routing and segment-level routing, and developing efficient decoding methods for Lory.

Ethics Statement

This paper presents a new approach for building large language models. We would like to note that, similar to existing language models, the language models trained with our approach may have the same potential societal consequences. For example, language models can produce factually inaccurate outputs (e.g., (Min et al., 2023)), facing the risk of spreading misinformation; malicious users can extract training data that is used to train language models (Carlini et al., 2021), causing privacy and license problems. We acknowledge these potential negative consequences and caution those who use our approach to build powerful language models.

References

- Mikel Artetxe, Shruti Bhosale, Naman Goyal, Todor Mihaylov, Myle Ott, Sam Shleifer, Xi Victoria Lin, Jingfei Du, Srinivasan Iyer, Ramakanth Pasunuru, Giridharan Anantharaman, Xian Li, Shuohui Chen, Halil Akin, Mandeep Baines, Louis Martin, Xing Zhou, Punit Singh Koura, Brian O’Horo, Jeffrey Wang, Luke Zettlemoyer, Mona Diab, Zornitsa Kozareva, and Veselin Stoyanov. Efficient large scale language modeling with mixtures of experts. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 11699–11732, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.804. URL <https://aclanthology.org/2022.emnlp-main.804>.
- Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, 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.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pp. 2633–2650, 2021.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. Boolq: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint arXiv:1905.10044*, 2019.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- Damai Dai, Chengqi Deng, Chenggang Zhao, R. X. Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Y. Wu, Zhenda Xie, Y. K. Li, Panpan Huang, Fuli Luo, Chong Ruan, Zhifang Sui, and Wenfeng Liang. Deepseekmoe: Towards ultimate expert specialization in mixture-of-experts language models. *CoRR*, abs/2401.06066, 2024. URL <https://arxiv.org/abs/2401.06066>.
- William B. Dolan and Chris Brockett. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*, 2005. URL <https://aclanthology.org/I05-5002>.
- Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. Glam: Efficient scaling of language models with mixture-of-experts. In *International Conference on Machine Learning*, pp. 5547–5569. PMLR, 2022.

-
- William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *The Journal of Machine Learning Research*, 23(1):5232–5270, 2022.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. Unsupervised dense information retrieval with contrastive learning, 2021. URL <https://arxiv.org/abs/2112.09118>.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547, 2019.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*, 2017.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466, 2019.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. Race: Large-scale reading comprehension dataset from examinations. *arXiv preprint arXiv:1704.04683*, 2017.
- Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668*, 2020.
- Mike Lewis, Shruti Bhosale, Tim Dettmers, Naman Goyal, and Luke Zettlemoyer. Base layers: Simplifying training of large, sparse models. In *International Conference on Machine Learning*, pp. 6265–6274. PMLR, 2021.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Michael S Matena and Colin A Raffel. Merging models with fisher-weighted averaging. *Advances in Neural Information Processing Systems*, 35:17703–17716, 2022.
- Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. MetaICL: Learning to learn in context. In *NAACL-HLT*, 2022.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. FActScore: Fine-grained atomic evaluation of factual precision in long form text generation. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 12076–12100, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.741. URL <https://aclanthology.org/2023.emnlp-main.741>.
- Mohammed Muqeeth, Haokun Liu, and Colin Raffel. Soft merging of experts with adaptive routing. *arXiv preprint arXiv:2306.03745*, 2023.
- Joan Puigcerver, Carlos Riquelme, Basil Mustafa, and Neil Houlsby. From sparse to soft mixtures of experts. *arXiv preprint arXiv:2308.00951*, 2023.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.

-
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–16. IEEE, 2020.
- Stephen Roller, Sainbayar Sukhbaatar, Jason Weston, et al. Hash layers for large sparse models. *Advances in Neural Information Processing Systems*, 34:17555–17566, 2021.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan LeBras, and Yejin Choi. Socialiqa: Commonsense reasoning about social interactions. *arXiv preprint arXiv:1904.09728*, 2019.
- Noam Shazeer. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*, 2020.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017.
- Weijia Shi, Sewon Min, Maria Lomeli, Chunting Zhou, Margaret Li, Victoria Lin, Noah A Smith, Luke Zettlemoyer, Scott Yih, and Mike Lewis. In-context pretraining: Language modeling beyond document boundaries. *arXiv preprint arXiv:2310.10638*, 2023.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pp. 1631–1642, 2013.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. Fever: a large-scale dataset for fact extraction and verification. *arXiv preprint arXiv:1803.05355*, 2018.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023a.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023b.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *International Conference on Learning Representations (ICLR)*, 2018.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Edouard Grave. Ccnet: Extracting high quality monolingual datasets from web crawl data. *arXiv preprint arXiv:1911.00359*, 2019.
- Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International conference on machine learning*, pp. 23965–23998. PMLR, 2022.
- Fuzhao Xue, Zian Zheng, Yao Fu, Jinjie Ni, Zangwei Zheng, Wangchunshu Zhou, and Yang You. Openmoe: An early effort on open mixture-of-experts language models. *arXiv preprint arXiv:2402.01739*, 2024.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.

-
- Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28, 2015.
- Zexuan Zhong, Tao Lei, and Danqi Chen. Training language models with memory augmentation. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 5657–5673, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.382. URL <https://aclanthology.org/2022.emnlp-main.382>.
- Yanqi Zhou, Tao Lei, Hanxiao Liu, Nan Du, Yanping Huang, Vincent Zhao, Andrew M Dai, Quoc V Le, James Laudon, et al. Mixture-of-experts with expert choice routing. *Advances in Neural Information Processing Systems*, 35:7103–7114, 2022.
- Barret Zoph, Irwan Bello, Sameer Kumar, Nan Du, Yanping Huang, Jeff Dean, Noam Shazeer, and William Fedus. St-moe: Designing stable and transferable sparse expert models. *arXiv preprint arXiv:2202.08906*, 2022.

A Pseudocode of Causal Segment Routing

We include a pytorch-style pseudocode of our proposed *causal segment routing* strategy in Algorithm 1.

Algorithm 1 Pseudocode of causal segment routing.

```
# B: batch size (number of training instances)
# L: length of each training instance
# d: hidden dimension
# E: number of experts
# T: length of each segment
# R: routing network (input: hidden rep, output: routing weights)

input x # x: input tensor (BxLxd)

N = L // T # number of segments in each sample
seg_x = x.view(B*N, T, d) # split x into segments

# representation of each segment (BNxE)
repr = mean(seg_x, dim=1)

# routing results (not causal) (BNxE)
e = softmax(R(repr), dim=-1)

# routing results for the first segment
e_first = e.view(B, N, E)[: , 0]

# make causal routing results (shift 1)
e = roll(e, 1) # shift by 1

# set routing results of the first segment
e = e.view(B, N, E) # back to the instance view
e[:, 0] = stop_grad(e_first) # assign w/ stop gradient
e = e.view(B*N, E)

# MoE FFN forward with expert weights e
seg_y = moe_ffn(seg_x, e) # seg_y: B*N x T x d

# back to the instance view
y = seg.y.view(B, L, d)

return y
```

moe_ffn: compute the merged expert and process the input (equation 2).

B Computational Overhead of Routing and Merging

Here we investigate the computational overhead of our MoE models compared to the dense counterpart. We consider an MoE layer and an input tensor x consisting of L tokens and d dimensions: $x : L \times d$. We assume that the model uses SwiGLU as the activation function in FFNs and it up-projects the input x to d' -dimensional activations in FFNs. In this case, processing the input on an FFN requires roughly $6 \times L \times d \times d'$ FLOPs (there are two up projections and one down projections in SwiGLU-based FFNs). The overhead of soft-routing MoE comes mainly from the merging operation. Suppose that there are E experts and that the model makes a routing decision for every segment of T tokens (equivalently, there are L/T routing decisions). Each merging operation on E experts takes $6 \times E \times d \times d'$ FLOPs (we compute three merged matrices). Therefore, the total overhead will be $\frac{L}{T} \times 6 \times E \times d \times d'$ FLOPs. This indicates that compared to a dense FFN layer, an MoE layer with E experts requires $\frac{E}{T}$ more FLOPs, compared to the dense counterpart. In our experiments, we set $T = 256$; this suggests that using $E = 8$ experts introduces 3.1% more computations and using $E = 32$ experts introduces 12.5% more computations at the FFN/MoE layers. It is worth noting that the computations from FFN layers are only a subset of the full model computations, so 3.1% is an overhead upperbound when measuring on full models. In our experiments, our most straightforward implementation leads to a 15% or 28% slowdown of training efficiency when using 8 or 32 experts (Table 2). We discuss the training strategy for scaling up model sizes in Appendix H.1 and leave a better-optimized implementation for future work.

Model	Throughput (tokens/s/gpu)
0.3B	29,000
0.3B/8E	24,500
0.3B/16E	22,900
0.3B/32E	20,800

Table 2: Training throughput (tokens/s/gpu) of our MoE models and the dense counterpart. Our implementation is based on data parallelism with the ZeRO optimization (Rajbhandari et al., 2020).

C Details of Similarity-based Data Batching

We adapt the pipeline of in-context pre-training (Shi et al., 2023) in our approach. Given a set of documents \mathcal{D} , for each document $d \in \mathcal{D}$, we first use Contriever (Izacard et al., 2021) to retrieve top- k most similar documents $N(d)$. The similarity between the document d_i and d_j is defined as the cosine similarity of their Contriever embeddings, i.e., $\text{sim}(d_i, d_j) = \cos(C(d_i), C(d_j))$, where C denotes the Contriever encoder model. We implement an efficient approximate nearest-neighbors search based on the FAISS library (Johnson et al., 2019). Then, we sort all the documents according to the similarity and construct training instances by batch consecutive documents. We use the same greedy algorithm as Shi et al. (2023). We start from a single document and repeatedly add the document that has the highest similarity value and has not been added to the list; we restart the process with a new document if all documents that are connected to the last document of the list are selected. We repeat this process until there are no documents left.

Model	n_{params}	N	D	n_{head}
0.3B	0.3B			
0.3B/8E	1.8B	24	1024	16
0.3B/16E	3.5B			
0.3B/32E	6.8B			
1.5B	1.5B			
1.5B/8E	7.8B	48	1536	24
1.5B/16E	15.0B			
1.5B/32E	29.5B			

Table 3: Model architectures and sizes used in our experiments. For MoE models, we replace each FFN layers with a MoE layer. kE (e.g., “16E” in “0.3B/16E”) represents the architecture in which each FFN layer is replaced with a MoE layer of k experts. N : number of layers; D : hidden dimension of the model; n_{head} : number of attention heads.

D Model Configurations

In our experiments, we employ Lory to train decoder-only models which consists of effective parameters of 0.3B and 1.5B. For each FFN layer in the Transformer model, we replace it with MoE layers with E ($E \in \{8, 16, 32\}$) experts with exactly the same architecture. Table 3 shows the configurations of model architectures.

E Experiments on 7B models

Experimental Setups. We conduct experiments on a 7B architecture. Table 4 shows the configuration of the model architectures. We train a dense 7B model and a 7B/4E MoE model. For the 7B models, we follow LLaMA2 (Touvron et al., 2023b) and use a combination of several corpora as the training set. We down-sample the full training set to a subset of 200B tokens for 7B models. Due to limited resources, we only conduct experiments on randomly batched training data for 7B models and do not apply the similarity-based batching method.

Model	n_{params}	N	D	n_{head}
7B	7B	32	4096	32
7B/4E	19.7B			

Table 4: Model architectures and sizes used in our 7B experiments. For MoE models, we replace each FFN layers with a MoE layer. N : number of layers; D : hidden dimension of the model; n_{head} : number of attention heads.

Model	PIQA	SIQA	BoolQ	HellaSwag
1.5B/8E (prompt)	72.1	45.2	62.0	43.6
1.5B/8E (segment)	72.1	45.6	60.2	43.9
1.5B/16E (prompt)	71.3	45.0	56.0	43.7
1.5B/16E (segment)	72.9	45.4	55.2	43.6

Model	Wino	NQ	TQA	Avg
1.5B/8E (prompt)	63.7	7.3	24.2	45.4
1.5B/8E (segment)	61.8	7.3	24.4	45.1
1.5B/16E (prompt)	61.5	7.3	25.6	44.4
1.5B/16E (segment)	62.4	7.6	25.5	44.7

Table 5: Downstream performance of using different inference methods. We study two routing strategy for inference. *prompt*: we make the routing decision once on the entire input prompt; *segment*: we re-route and get new merged FFNs every segment.

Language Modeling Results. We show the training loss curves in Figure 7 and the perplexity on held-out evaluation sets in Table 6. We find that compared to the 0.3B and 1.5B models (see Section 4), the improvement of the 7B/4E model is less significant. We think it is because (1) the similarity-based batching method is not applied in this case, making the experts under-utilized; (2) we only use four experts in the MoE model. We leave the experiments with the similarity-based batching method on MoE models with more experts as future work.

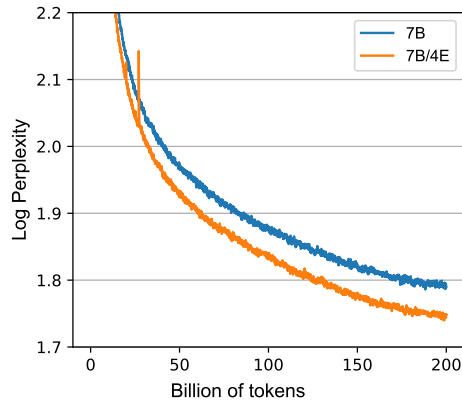


Figure 7: Training curves (log perplexity) of the 7B dense model and the 7B/4E MoE model. Note that when training the 7B/4E model, we do not apply the similarity-based batching method.

Performance on Downstream Tasks.

Table 7 shows the performance of the models on downstream tasks. We find that although the similarity-based batching method is not used when training the 7B/4E model, it still achieves clearly better results on various tasks compared to the dense 7B model. This further suggests the effectiveness of our causal routing strategy.

Model	arXiv	Books	Wiki	C4	Python
7B	2.3	9.1	5.9	8.0	2.3
7B/4E	2.2	8.7	5.7	7.7	2.2

Table 6: Perplexity of trained models on different evaluation sets (arXiv, Books, Wikipedia, C4, and Python). Note that when training the 7B/4E model, we do not apply the similarity-based batching method.

Model	Commonsense Reasoning					Reading Comprehension			
	PIQA	SIQA	BoolQ	HellaSwag	WinoGrande	RACE-m	RACE-h	ARC-e	ARC-c
7B	76.9	50.2	65.2	52.6	66.2	55.3	40.5	73.0	38.5
7B/4E	77.7	50.1	67.6	54.8	67.3	57.0	41.3	73.5	39.6

Model	Closed-book QA		Text Classification					Avg	
	NQ	TQA	AGNews	Amazon	SST-2	Yelp	Fever		MRPC
7B	17.3	42.5	80.6	94.3	92.7	98.3	53.7	67.0	62.5
7B/4E	18.8	44.5	81.7	95.7	93.1	96.7	57.7	69.7	63.9

Table 7: We compare the 7B/4E MoE models trained with our routing strategy (without using the similarity-based batching method) with the parameter-matched dense models on downstream tasks, including commonsense reasoning, reading comprehension, closed-book QA, and text classification.

F Expert Specialization: Full Results of Routing Weights

Figure 12 shows the routing weights at all layers of the 0.3B/8E model. It clearly shows that our MoE models are able to learn domain-level specialization.

G More Analysis and Ablation Studies

G.1 Held-out Evaluation of Lory and EC models

Table 8 shows the perplexity of the models on held-out evaluation sets. We find that the token-level EC model performs better than our model on the C4 evaluation data, which is likely on the most similar distribution to the training set (Commoncrawl); on arXiv, Books, and Wikipedia, EC performs similarly or slightly worse than our model. Interestingly, on the Python evaluation set, our model achieves a much better perplexity compared to the token-level EC model (12.5 vs 13.6). We suspect that segment-level routing may be particularly effective at handling out-of-domain evaluation data (we assume that there is only a very small part of Python code in Commoncrawl). Our analysis in Section 5.4 shows that segment-level routing models are indeed able to learn experts that are specialized in specific domains (e.g., Python code), potentially helping models achieve high performance in less frequent domains.

Model	arXiv	Books	Wiki	C4	Python
0.3B/8E (Lory)	7.4	16.0	9.2	13.3	12.5
0.3B/8E (EC, seg-level)	7.9	17.6	10.5	14.1	14.1
0.3B/8E (EC, tok-level)	7.5	17.0	9.2	12.8	13.6

Table 8: Perplexity of our trained MoE model and EC models on evaluation sets. We instantiate EC methods with our segment-level routing and the original token-level routing.

G.2 Expert Utilization

Here, we investigate the expert utilization of our MoE models. We define an expert that is activated for a given input if the routing weight is larger than $\frac{2}{E}$ (e.g., larger than 25% if there are 8 experts each layer). In Figure 8, we plot the number of experts that are activated at least once among 10 training steps (consisting of 10M tokens and 40K segments) when training 1.5B MoE models. We see that after the warmup phase at the beginning, the expert utilization quickly increases. 1.5B/8E and 1.5B/16E models have quickly utilized most of




Figure 8: We show how many experts are actively utilized at least once every 10 training steps. We define an expert is activated if the weight is larger than $\frac{2}{E}$, where E denotes the number of experts at each MoE layer. Our models have 48 layers; therefore, the 1.5B/8E, 1.5B/16E, 1.5B/32E models have 384, 768, 1536 experts in total, respectively.

Model	PIQA	SIQA	BoolQ	HellaSwag
1.5B/8E (prompt)	72.1	45.2	62.0	43.6
1.5B/8E (segment)	72.1	45.6	60.2	43.9
1.5B/16E (prompt)	71.3	45.0	56.0	43.7
1.5B/16E (segment)	72.9	45.4	55.2	43.6
Model	Wino	NQ	TQA	Avg
1.5B/8E (prompt)	63.7	7.3	24.2	45.4
1.5B/8E (segment)	61.8	7.3	24.4	45.1
1.5B/16E (prompt)	61.5	7.3	25.6	44.4
1.5B/16E (segment)	62.4	7.6	25.5	44.7

Table 9: Downstream performance of using different inference methods. We study two routing strategy for inference. *prompt*: we make the routing decision once on the entire input prompt; *segment*: we re-route and get new merged FFNs every segment.

the experts; while the expert utilization of the 1.5B/32E model continues to increase until the end of the training and it is able to activate about 900 experts among 1536 experts at the end of training. This indicates that our approach is able to prevent the MoE models from collapsing to dense models and achieves high expert utilization. However, when training with a large number of experts, achieving high expert utilization is more challenging.

G.3 Inference Methods

During inference of downstream tasks, by default, we take the task input prompt as the input of the routers in each layer and make the routing decision once. This inference method enables the decoding process to be simple and achieves low latency, since after encoding and routing the input, we do not need to use the routers again – the rest generation can be run in a (merged) dense model. As such an inference method introduces a train-test discrepancy, we study the method that routes every segment as we do during training. As shown in Table 9, routing the input once or routing each segment does not make substantial differences in the downstream tasks we evaluate. Due to simplicity and efficiency, we use the entire prompt as the routing input and perform routing only once.

In addition, we discuss the possibility of converting Lory to sparse models for more efficient inference in Appendix H.2.

G.4 Warmup Training

At the beginning of training (i.e., the first 5% training steps), we train a dense LM with the same configuration before training the MoE model. We initialize the MoE layers by duplicating the FFN layers of the warmup trained model. We find that this warmup phase

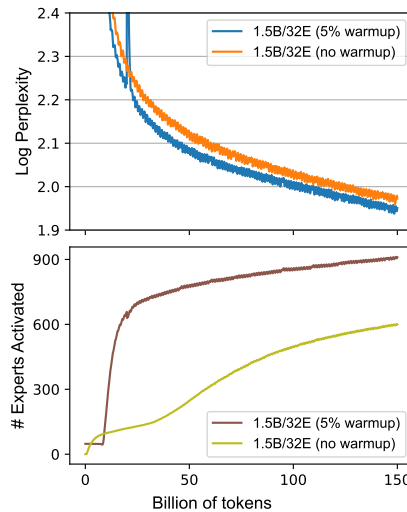


Figure 9: Training curves and expert utilization of employing a warmup phrase or not. We find without a warmup phrase, training leads to a worse MoE model (top) and worse expert utilization (bottom).

is crucial for achieving high expert utilization especially when there is a large number of experts. Figure 9 visualizes the training loss curves and expert utilization of the 1.5B/32E model (with or without warmup training). As shown in the figure, without the warm-up phrase, the model achieves worse performance and much fewer experts are utilized.

H Discussion

H.1 Parallelism for Scaling

In our experiments, we use data parallelism with ZeRO optimization (Rajbhandari et al., 2020) to reduce memory usage. Data parallelism either saves a copy of the model parameters in each device or relies on fast communication between devices to synchronize parameters if using the ZeRO optimization. When we scale up the MoE models (e.g. model sizes > 100 B), this introduces a potential problem as the amount of MoE parameters increases dramatically, resulting in a bottleneck due to significant parameter communication.

To address this problem, expert parallelism has been utilized for sparsely activated MoE models, where experts are distributed across different devices and host independently (Lepikhin et al., 2020; Fedus et al., 2022). Unlike traditional MoE models that involve a top-k discrete routing network to dispatch inputs to a subset of experts, our soft-routing MoE models require access to all expert parameters for the computations. To overcome this challenge, we employ an expert-wise model parallelism strategy. This involves partitioning the MoE layers across multiple devices by sharding all experts along with their hidden dimension. Consequently, same dimensions from different experts are allocated to the same device. As Figure 10 shows, we shard all the experts along with the hidden dimension, and distribute the same shard of different experts to one device. The expert-wise model parallelism enables us to communicate only hidden activations instead of model parameters. This is much more efficient as the hidden activations do not scale with the number of experts.

Furthermore, for non-MoE components (e.g., attention layers), we can still employ data parallelism to avoid communication overhead. Figure 11 illustrates a training setup in which we employ a mixture of data parallelism and model parallelism for training at scale.

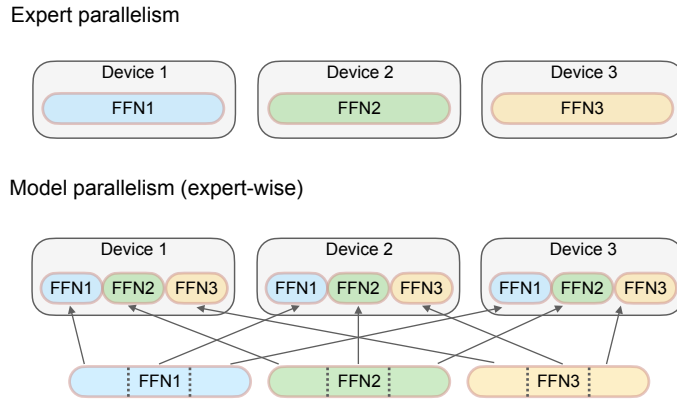


Figure 10: Comparison between expert parallelism and expert-wise model parallelism. In soft-routing MoE models, the access to all expert parameters is needed for the computations. We shard all the experts along with hidden dimension, and distribute them across different devices.




Figure 11: Illustration of using a mixed strategy of data parallelism and expert-wise model parallelism. For non-MoE modules (e.g., attention layers), we use the data parallelism; while for MoE layers, we employ the model parallelism to avoid significant parameter communication.

H.2 Converting to Sparse Models for Inference

In our MoE models, we merge FFN experts for each segment. This can lead to a large memory usage during inference when the model process a large inference batch, because the parameters of a merged FFN per batch are cached in the GPU memory. One possibility to alleviate this issue is to fine-tune the trained models with a hard-decision routing mechanism (e.g., top-k routing) after the pre-training stage. This method transitions the models with soft routers to ones with hard routers, significantly reducing the memory usage during inference. We leave further investigation on this direction as future work.



Figure 12: Averaged routing weights at all the layer of the 0.3B/8E model on different domains (Books, arXiv, Python, Wikipedia). We observe that the experts in our MoE models learn clear domain-level specialization.