

NO NEED TO TALK: ASYNCHRONOUS MIXTURE OF LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

We introduce SMALLTALK LM, an innovative method for training a mixture of language models in an almost asynchronous manner. Each model of the mixture specializes in distinct parts of the data distribution, without the need of high-bandwidth communication between the nodes training each model. At inference, a lightweight router directs a given sequence to a single expert, according to a short prefix. This inference scheme naturally uses a fraction of the parameters from the overall mixture model. Our experiments on language modeling demonstrate that SMALLTALK LM achieves significantly lower perplexity than dense model baselines for the same total training FLOPs and an almost identical inference cost. Finally, in our downstream evaluations we outperform the dense baseline on 75% of the tasks.

1 INTRODUCTION

Recent research has demonstrated that scaling large language models (LLMs) by increasing model capacity and expanding training data consistently leads to significant performance improvements on a wide range of downstream tasks (Kaplan et al., 2020; Brown et al., 2020; Henighan et al., 2020; Hoffmann et al., 2022; Dubey et al., 2024). Scaling introduces substantial operating and engineering costs for both inference and training. In general, training is achieved on a large number of nodes via synchronous gradient descent techniques, which relies on high-bandwidth communication to scale. Inference of large models may require multiple compute nodes to distribute the model, which relies on low-latency communication. In both cases, state-of-the-art interconnect hardware is critical, and careful engineering is required to scale, and maintain a large compute cluster. While mainstream machine learning frameworks have eased the engineering work on the scaling side, access to a large number of well interconnected nodes remains a privilege in the machine learning community.

In this paper, we explore strategies to mitigate the communication cost of large language models, both at training and inference, while keeping the inference efficient. We show that efficient training and inference can be achieved without relying on fast interconnects, and without compromising model performance, both in terms of perplexity or downstream task accuracy.

In recent studies aimed at mitigating reliance on high-bandwidth interconnects, researchers have developed algorithms that reduce the need for frequent or comprehensive gradient synchronizations. Such techniques include asynchronous training (Douillard et al., 2023; Aji & Heafield, 2017; Zhang et al., 2015; Liu et al., 2024) and gradient compression (Lin et al., 2017; Dettmers, 2015) which effectively decrease communication overhead. By performing updates less frequently or by communicating less data, these methods sustain high training throughput while diminishing dependence on high-speed interconnects. However, these algorithms still require some level of gradient synchronization, and resulting models often under perform (in terms of perplexity) compared to training approaches synchronizing at every step (Diskin et al., 2021).

Regarding efficient inference, a number of sparse parameter activation techniques have recently become popular (Shazeer et al., 2017; Fedus et al., 2022; Artetxe et al., 2022; Lewis et al., 2021; Du et al., 2022), in particular the Switch Transformer mixture of experts (MoE). These approaches are effective at reducing significantly the number of active model parameters at inference time. Yet, they require routing decisions to be made for each token, demanding rapid interconnections and requiring all parameters to be accessible in RAM.

To achieve both low-bandwidth training and efficient inference, several prior works have explored asynchronously training mixtures of models (Gross et al., 2017; Li et al., 2022; Gururangan et al., 2023). In these methods, the input space is clustered using unsupervised techniques like K-Means, with each model in the mixture trained on an assigned cluster. These approaches reduce communication overhead and have demonstrated improvements in model perplexity, while maintaining similar inference costs to baseline models (Gross et al., 2017; Gururangan et al., 2023). However, in the context of language processing, their practical applicability on standard downstream tasks with relatively short input prefixes remains a subject of debate, as they rely on full corpus clustering or access to metadata (Li et al., 2022; Gururangan et al., 2023).

In this paper, we introduce SMALLTALK LM, an asynchronous mixture of language models, which combines the advantage of asynchronous LM training methods (significantly reducing interconnect requirements), and sparse activation methods (it naturally uses only a subset of the overall architecture parameters during inference). SMALLTALK LM achieves better perplexity, and better accuracy on a majority of downstream tasks, compared to a regular dense language model trained with the same amount of FLOPs. Our main contributions are as follows:

- We introduce an algorithm that trains a mixture of independent language models, significantly reducing bandwidth requirements during training compared to regular distributed training methods. Our approach leverages a lightweight router (which accounts for less than 1.5% of the size of each expert) to efficiently route sequences. This router determines the most suitable expert for a given sequence based on a short prefix (see § 2).
- We empirically demonstrate that our method achieves significantly lower perplexity than a dense baseline for near-identical training and inference FLOPs and identical training data volume (see Fig. 2). Furthermore, we show that for a constant mixture model size, the improvement in perplexity increases with the number of experts.
- We evaluate our method on a set of downstream tasks, showing that it achieves better or similar accuracy compared to a dense baseline on 75% of the tasks (see § 3.3 and App. B).

2 METHOD

In this section, we formalize our proposed method: SMALLTALK LM. In § 2.1, we introduce general language modeling background, and present a mixture of experts framework in this context. Subsequently, in § 2.2, we show how we can use independent language models to perform routing to different experts in the mixture. We also present the training procedure and implementation details of our method. Finally, we discuss its marginal computational overhead and minimal distributed communication costs.

2.1 BACKGROUND

Language modeling Let \mathcal{V} denote a fixed vocabulary, and consider a sequence of tokens $\mathbf{x} = (x_1, x_2, \dots, x_S)$, where each $x_s \in \mathcal{V}$ for $s = 1, 2, \dots, S$. Language modeling is about modeling the data distribution $p(\mathbf{x})$, which in practice is cast as training a language model $p(x_{s+1} | \mathbf{x}_{1:s}; \theta)$ generating a token x_{s+1} given a context sequence $\mathbf{x}_{1:s} = (x_1, x_2, \dots, x_s)$. Typically, the model parameterized by θ is a neural network, and parameters are obtained by minimizing the negative log-likelihood:

$$\mathcal{L}(\mathbf{x}; \theta) = -p(\mathbf{x}; \theta) = -\sum_{s=1}^{S-1} \log p(x_{s+1} | \mathbf{x}_{1:s}; \theta). \quad (1)$$

Mixture of Experts There exist many variants of the mixture of experts model (Jacobs et al., 1991; Jordan & Jacobs, 1994; Tresp, 2000; Collobert et al., 2001; Shahbaba & Neal, 2009). Following the original formulation and applying it to language modeling, we factorize the conditional distribution for predicting a given token by summing over a set of E experts:

$$\mathcal{L}(\mathbf{x}; \theta) = -\sum_{s=1}^{S-1} \log \sum_{e=1}^E p(x_{s+1} | \mathbf{x}_{1:s}, e; \theta^e) p(e | \mathbf{x}_{1:s}; \theta^r), \quad (2)$$

where each expert $p(x_{s+1} | \mathbf{x}_{1:s}, e; \theta^e)$ is intended to focus on different parts of the data distribution, and the router $p(e | \mathbf{x}_{1:s}; \theta^r)$ assigns weights to the expert predictions. Each expert e has its own set of independent parameters θ^e , and the router is parameterized by θ^r .

While this type of mixture of experts partitions the input space across different models, all experts are still involved when computing the conditional likelihood in Equation (2), both at training and inference time. Hard mixtures of experts circumvent this issue by hard-assigning a sequence $\mathbf{x}_{1:s}$ to a *single* expert e^* , minimizing an upper bound of Equation (2):

$$\mathcal{L}(\mathbf{x}; \theta) \leq - \sum_{s=1}^{S-1} \log p(x_{s+1} | \mathbf{x}_{1:s}, e^*; \theta^{e^*}) p(e^* | \mathbf{x}_{1:s}; \theta^r), \quad (3)$$

The way e^* is chosen depends of the goal of the approach. For example, in [Jacobs et al. \(1991\)](#), e^* is sampled at training time, according to the distribution $p(e | \mathbf{x}_{1:s}; \theta^r)$. In other works, like [Collobert et al. \(2001\)](#), it is chosen as the most relevant expert (that is $e^* = \arg \max_e p(e | \mathbf{x}_{1:s}; \theta^r)$), both at training and inference. Regarding the training procedure, mixture models with both soft and hard assignments are typically trained with an Expectation-Minimization (EM) scheme, alternating between experts and router optimization.

2.2 SMALLTALK LM

Routing With Independent Language Models SMALLTALK LM can be viewed as an instance of hard mixture of experts, where a sequence is routed to the most relevant expert e^* , according to the router posterior $p(e | \mathbf{x}_{1:s}; \theta^r)$, both at training and inference. However, instead of being a monolithic neural network model, the router is implemented as a language model per expert. More precisely, with Bayes’ rule we have:

$$e^* = \arg \max_e p(e | \mathbf{x}_{1:s}; \theta^r) \quad (4)$$

$$= \arg \max_e \frac{p(\mathbf{x}_{1:s} | e; \theta^r) p(e)}{\sum_{i=1}^E p(\mathbf{x}_{1:s} | i; \theta^r) p(i)} \quad (5)$$

$$= \arg \max_e p(\mathbf{x}_{1:s} | e; \theta^r) \quad (6)$$

$$= \arg \max_e p(\mathbf{x}_{1:s} | \theta^{r,e}), \quad (7)$$

where the likelihood $p(\mathbf{x}_{1:s} | \theta^{r,e})$ defines a routing language model attributed to the e -th expert with *independent* parameters $\theta^{r,e}$, and expert priors $p(e)$ are supposed to be uniform. So far, in our definition, there is nothing that prevents $\theta^{r,e} = \theta^e$, namely to use the experts themselves to implement the router. Interestingly, in that case selecting e^* as above is intuitive because it amounts to selecting the best expert for a given sequence (in terms of log-likelihood).

Deriving a Practical Model In order for the above formulation to result in a model that is useful in practice, we have to solve the following main challenges. Firstly, we need to be able to route using a small prefix of the full sequence since in most real world applications we do not have access to a full sequence (e.g. in question answering we do not have the answer). Secondly, using the experts to implement the router results in a large computational overhead since we need to evaluate every single one of them, largely negating the benefits of the hard mixture of experts.

We solve both of the above problems with the following two key ideas:

1. We use a *short prefix* of length M to compute the score for each expert as follows:

$$p(e^* | \mathbf{x}_{1:s}) \approx p(e^* | \mathbf{x}_{1:M}). \quad (8)$$

This enables us to use SMALLTALK LM in real-world applications where only a very small prefix may be available. We show in § 3.4 that our method outperforms the dense baseline even when using as little as 32 tokens for routing.

2. We implement the router with E language models that are *orders of magnitude smaller* than the experts. This results in marginal computational overhead, as little as 1% during training and 3% during inference (§ 3.2).

Using the log-likelihood of a prefix to route a sequence is an intuitive solution since a model that is good at predicting the beginning of a sequence is likely to be good at the full sequence. Another interpretation is that the prefix can be used to perform some sort of thematic identification of the whole sequence that is then used to route to the appropriate expert.

On the other hand, the efficacy of smaller language models as routers is less intuitive. Even though one can expect a large model to be able to perform well on sequences that are beyond the capabilities of smaller models, the task of routing boils down to identifying whether a sequence is similar to the training data or not. We show that even really small models (4.4M parameters) can perform that task very well and result in no discernible difference in performance compared to larger routers (§ 3.4).

Algorithm 1 SMALLTALK LM training.

```

1: Train the routers
2:  $\mathbf{X} \leftarrow N$  new sequences from the dataset
3:  $\mathbf{X}_{1:E} = \text{random\_assignments}(\mathbf{X})$ 
4: for  $i = 1 \dots T$  do
5:   for  $e = 1 \dots E$  do
6:      $\theta^{r,e} \approx \arg \min_{\theta^{r,e}} \mathcal{L}(\mathbf{X}_e; \theta^{r,e})$  # Optimize Equation 9 with SGD for the  $e$ -th router
7:   end for
8:    $\mathbf{X} \leftarrow N$  new sequences from the dataset
9:    $\mathbf{X}_{1:E} = \text{balanced\_assignments}(\mathbf{X}, \theta^r)$  # Segment the data according to Equation 4
10: end for
11: Train the experts
12:  $\mathbf{X} \leftarrow M$  new sequences from the dataset comprising the total number of training tokens
13:  $\mathbf{X}_{1:E} = \text{balanced\_assignments}(\mathbf{X}, \theta^r)$ 
14: for  $e = 1 \dots E$  do
15:    $\theta^e \approx \arg \min_{\theta^e} \mathcal{L}(\mathbf{X}_e; \theta^e)$  # Optimize Equation 1 with SGD for the  $e$ -th expert
16: end for

```

Training Procedure We train our proposed hard mixture of experts with EM. Namely we alternate between a step of optimizing equation 1 and a step of assignments according to equation 4. In our framework, assignments e^* purposefully do not depend on the experts, and the routers are independent language models. This allows us to split the training in two stages:

1. We train the routers using EM without using the experts at all, namely we minimize

$$\mathcal{L}(\mathbf{x}, \theta^r) = -\log p(e^* | \mathbf{x}_{1:M}; \theta^r) = -\sum_{s=1}^{M-1} \log p(x_{s+1} | \mathbf{x}_{1:s}; \theta^{r,e^*}), \quad (9)$$

in the likelihood maximization step and perform the assignments using equation 4 in the expectation step.

2. We train the experts as E independent language models on disjoint dataset segments selected by the trained routers.

This procedure is described in pseudocode in Algorithm 1. In the specific case where the experts and the routers are the same language models ($\theta^e = \theta^{r,e}$) our training algorithm comes down to a regular EM algorithm where the sequences are routed to the best performing expert. In that light, we can think of each router as lightweight approximation of the corresponding expert. See § 3.4 for experiments where the experts themselves are used for routing.

Balancing the Assignments to Experts A common hurdle to overcome when training mixture of experts models is to ensure that the modeling load is shared as evenly as possible across all experts, or in other words that the model is utilizing the capacity of all experts efficiently instead of relying on a few good ones. Common ways of tackling this problem are either introducing losses that penalize uneven assignments or introducing noise into the assignments namely by sometimes performing “sub-optimal” assignments to avoid local-minima. These techniques are used in all mixture models since their introduction in [Jacobs et al. \(1991\)](#) until more recently in [Fedus et al. \(2022\)](#).




Figure 1: **Balanced assignments.** In this scenario we have 3 experts (columns with different colors) and 3 sequences to assign (rows) under the constraint that each expert should get 1 sequence. In (a) we assign sequentially each row, by negative log-likelihood, leading to sub-optimal assignment because the first expert is full when we try to assign the last row. On the other hand, in (b), we first sort wrt to the minimum log-likelihood which results in optimal assignments.

In our setup, we can solve this problem by ensuring that each expert is assigned equally sized distinct portions of the dataset. However, a challenge arises when assigning data sequentially based solely on Equation 4. In such cases, an expert might already have more than its fair share of the dataset assigned yet we continue to find sequences for which it has the lowest loss among all experts, as illustrated in Figure 1a. To overcome this issue, during training we perform the assignments considering the whole set of sequences rather than one sequence at a time. In particular, we sort the sequences based on $-\max_e \log p(\mathbf{x}_{1:M} | e; \theta^r)$ and perform the assignments in that order. This ensures that sequences with high likelihood will be assigned first and avoids the edge-case where a sequence with low-loss cannot be assigned to the appropriate expert because it is at capacity. In Algorithm 1 we refer to this procedure as *balanced assignments* and it is demonstrated in Figure 1b.

During inference, no balancing is performed, and the expert is selected solely based on equation 4.

Computational and Communication Cost It is important to note that the computational cost of training and inference with the routers is significantly lower than that of training and inference with the experts. This difference arises due to the routers’ much smaller size, smaller batch sizes, and the fewer tokens they train on (see Appendix A.3 for further details). As a result, the additional computational overhead from training and inference with the routers is negligible.

Regarding communication requirements for a distributed implementation, in Algorithm 1 all operations except for the balanced assignments are completely independent and embarrassingly parallelizable without any communication. Namely, each expert or each router can be trained on its own on a different node or set of nodes. The assignments, however, require that every node has access to the scores of every router in order to perform the assignment according to equation 4. This requires a small amount of communication as each node needs to share 1 score per sequence which results in 2MB per 1M sequences if we represent the scores in 16-bit floats. A detailed analysis is provided in the Appendix A.4.

3 EXPERIMENTS

In this section, we experimentally analyze the performance of our proposed method. First, in § 3.2 we benchmark our approach against a dense LLM trained using regular distributed training techniques. Second, in § 3.3, we evaluate our method on downstream tasks, demonstrating that it achieves performance comparable to baseline models at the same level of perplexity, improving upon a dense model given the total training FLOPs spent. Finally, in § 3.4, we analyse the core component of our method, the routing. We show that the performance of the mixture does not depend significantly on the size of the router. Furthermore, we show that our approach outperforms less complex routing methods, such as the TF-IDF encoding with K-Means clustering proposed by Gururangan et al. (2023). Moreover, we demonstrate that the perplexity gains are not driven by a few dominant experts but result from each expert contributing positively to the overall performance.

270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323

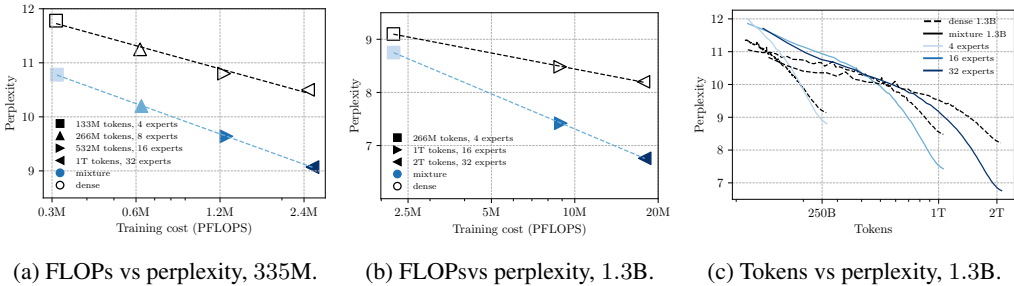


Figure 2: **Better perplexity for the same price.** Test perplexity comparison between our approach and the dense baseline, as a function of training cost measured in PFLOPs. In (a), we report results for models with 335M parameters using 4, 8, 16, and 32 experts and in (b) for models with 1.3B parameters using 4, 16, and 32 experts. In addition, (c) shows the perplexity comparison between our approach and the dense baseline, plotted against the cumulative number of tokens processed throughout the training for the 1.3B parameter models. We observe that our method significantly outperforms the baseline across all experimental configurations. Notably, our 335M parameter model with 32 experts achieves a perplexity of 9.07, outperforming the 1.3B dense baseline’s perplexity of 9.1. This improvement is achieved with a training budget of 2.5×10^{21} FLOPs, which is comparable to the baseline’s 2.2×10^{21} FLOPs, while requiring **three times less** computational cost during inference (0.87×10^{12} FLOPs compared to 2.81×10^{12} FLOPs). See § 3.2 and App. A for a detailed description of our experimental setup.

3.1 EXPERIMENTAL SETUP

We use the RedPajama-V2 dataset (Computer, 2023), which is a large-scale collection of text data designed for training language models. The dataset is built from the ground up based on publicly available web data, consisting of 84 crawls provided by Common Crawl (2024).

Both the experts and routers utilize a Transformer architecture (Vaswani, 2017) with rotary positional encoding (Su et al., 2024). All experiments use sequences of 1,024 tokens. We train the models using the AdamW optimizer (los) with $\beta_1 = 0.9$, $\beta_2 = 0.99$, and a weight decay of 0.1. Gradient clipping is applied with a maximum norm of 0.1.

For the experts, we employ a learning rate schedule featuring a linear warm-up to 5×10^{-4} over the first 3,000 steps, followed by a cosine decay for the remainder of the training period. We experiment with two model sizes: a 335M-parameter model with 4, 8, 16, and 32 experts and a 1.3B-parameter model with 4, 16, and 32 experts. Tables 1 and 2 describe the architectures and the training parameters in more detail.

For the routers, we use models with 4.4M parameters. The routers are trained for 128,000 steps using a constant learning rate of 1×10^{-4} , following a linear warm-up over the first 1,000 steps. Routers are trained with a batch size of 32 and perform all-to-all communication of the loss on the dataset chunk approximately every 45 million training tokens (see App. A.4). The length of the prefix used for routing is set to 256 tokens, which is 25% of the context length.

Comparison to the Dense Model To ensure a fair comparison to the dense baseline, we designed our experiments such that the computational cost during inference and the total computational cost during training are the same. We achieve this by comparing the performance of models where each expert has the same architecture and number of parameters as the dense model, resulting in the same inference cost. In addition, we train SMALLTALK LM for the same number of total tokens which means we spend the same number of total training FLOPs.

3.2 LANGUAGE MODELING RESULTS

Across all experimental configurations, our method *consistently outperforms the baseline* while utilizing the same volume of data, with only minor increases in computational costs. Specifically, for the 1.3B parameter model, our approach results in less than a 1% increase in training FLOPs and




Figure 3: **Downstream evaluation.** Accuracy with respect to perplexity on (a) ARC Challenge, (b) ARC Easy, (c) HellaSwag and (d) MMLU, for 1.3B parameter dense baselines trained on 266B, 1T and 2T tokens (empty symbols) and mixture models with 1.3B parameter experts and 4, 16 and 32 experts respectively (filled symbols). The models that have the same symbol shape have near identical training and inference FLOPs.

less than a 3% increase in inference FLOPs, while achieving performance gains of up to almost 18% in perplexity. Similarly, for the 335M parameter model, the cost increases are modest, with less than a 4% increase in training FLOPs and less than a 10% increase in inference FLOPs, alongside performance gains of up to nearly 14% in perplexity (see Appendix A.3 for details on FLOPs calculation). Figures 2a and 2b illustrate the perplexity on a held-out test set relative to the total training cost, measured in petaFLOPs, across different model sizes. Additionally, Figure 2c shows the test perplexity as a function of the number of training tokens for the 1.3B parameter model.

Moreover, it is interesting to compare our model with the smaller 335M parameter experts to the 1.3B dense baseline. When using 32 experts and roughly the same training budget in FLOPs, SMALLTALK LM achieves slightly better perplexity (9.07 vs 9.11) but requires **three times less** compute during inference.

Finally, we observed that for a constant model size, the improvement in perplexity increases with the number of experts. This suggests that our method more efficiently utilizes the training data to improve performance, likely due to the specialization of experts on different segments of the data distribution.

Communication Overhead Our approach performs minimal communication and does not require a fast interconnect. As detailed in § 2, we first train the routers which are then used to efficiently shard the dataset and each expert is trained completely independently. During their training, the routers communicate approximately 100 times with each transmission involving less than 6 megabytes per router resulting in truly minimal communication requirements. See A.4 for details about the calculation of these numbers.

3.3 DOWNSTREAM EVALUATION

To demonstrate the applicability of our approach in real-world scenarios, we performed zero-shot evaluation on a set of natural language processing (NLP) tasks, including ARC Challenge and ARC Easy (Clark et al., 2018), HellaSwag (Zellers et al., 2019), SciQ (Welbl et al., 2017), and MMLU (Hendrycks et al., 2020). Figure 3 shows both accuracy and perplexity for each task (see Appendix B for details regarding perplexity and accuracy computation).

Specifically, our largest configuration, a model with 32 1.3B parameter experts, demonstrates better performance on four out of five tasks, achieving gains of 3%, 2%, 3%, and 1% on ARC Challenge (3a), ARC Easy (3b), HellaSwag (3c), and MMLU (3d), respectively, under the same inference cost.

Moreover, on the MMLU benchmark, our approach either beats the baseline or achieves the same performance on 75% of the tasks (42 out of 56), as shown in Table 4.

3.4 ROUTING ANALYSIS

In this section we analyse the most critical component of SMALLTALK LM, the routing. In particular, we perform experiments to investigate the impact of the size of the router to the mixture model,

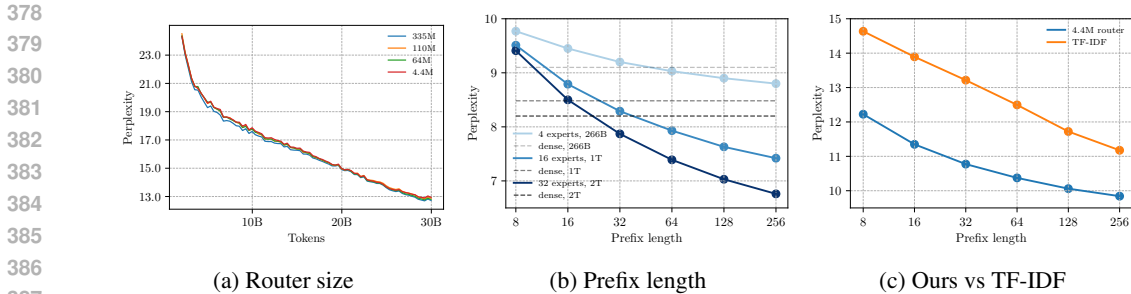


Figure 4: **Routing analysis.** (a) Test perplexity over training steps for different router sizes using a 335M parameter model with 4 experts. We compare routers of sizes 335M (where the model routes data for itself), 110M, 64M, and 4.4M parameters. (b) Test perplexity as a function of routing prefix length during inference for 1.3B parameter model with 4, 16 and 32 experts. We examine how reducing the prefix length \hat{M} used during inference affects performance when the data is partitioned during training using a prefix size $M \geq \hat{M}$. (c) Test perplexity over training steps for a 335M parameter model with 16 experts, comparing our proposed routing using TF-IDF document encoding followed by SVD projection and balanced K-Means clustering.

the impact of the size of the prefix and finally whether the EM algorithm presented in § 2 is critical or it could be replaced by simple content based clustering.

Impact of Router Size One of the most critical findings of our work is the fact that *the size of the model* used for routing *does not impact* the performance of the mixture (Figure 4a). This allows us to utilize small router models which reduce dramatically the computational cost of routing during training and inference without compromising the model’s performance.

In detail, we investigated the effect of router size on the performance of our method by experimenting with a 335M parameter model configured with 4 experts, using four different router sizes for data partitioning: a 335M parameter router (where the model routes data for itself, optimizing for the data it handles best), and routers with 110M, 64M, and 4.4M parameters. All models were trained with a batch size of 32 for 256,000 steps, using a routing prefix length of 256 tokens, following the same experimental setup as in § 3.2 (see Appendix A for more details about model architectures and training parameters). Figure 4a illustrates the test perplexity wrt the number of training tokens for these configurations. We observe that all router sizes perform practically identically.

Impact of Prefix Length An important hyperparameter for our mixture model is the length of the context used for routing to the expert model. In our experiments, we use a prefix length of 256 which may be too large for some real-world applications like conversational AI. However, we show in Figure 4b that SMALLTALK LM outperforms the dense baseline even when routing with shorter prefixes. In particular, we show that even with prefixes of just 32 tokens ($\frac{1}{8}$ th of the length during training) the model still outperforms the dense baseline.

Comparison With Routing Using TF-IDF An alternative approach might suggest that if a tiny LLM suffices for routing, simpler methods such as TF-IDF encoding combined with balanced K-Means clustering, as proposed by Gururangan et al. (2023), could be adequate. To test this hypothesis, we trained a 335M parameter model using the routing strategy outlined in (Gururangan et al., 2023): applying TF-IDF transformation to the text, followed by Singular Value Decomposition (SVD) projection into a low-dimensional space, and then clustering using balanced K-Means.

We then trained 16 experts on these clustered results. Our findings reveal that routing with TF-IDF encoding underperforms when the prefix is short, indicating limitations of this simple routing approach. As shown in Figure 4c, our proposed routing method significantly outperforms the TF-IDF-based routing. This demonstrates that our approach is more effective in capturing the nuances of the data distribution, even when using limited context for routing decisions.

Experts Do Specialize A potential concern is that the observed performance gains might be driven primarily by a few experts among the E , while the remaining contribute minimally or may even under perform on their assigned data. To address this concern, we conducted an analysis of each expert’s performance individually. Figures 5a, 5b, and 5c compare the perplexity achieved by our method and the dense baseline on the routed dataset segments for the 1.3B parameter model, utilizing mixtures of 4, 16, and 32 experts, respectively.

Our findings demonstrate that all experts contribute positively, with consistent perplexity improvements observed across all routed dataset segments. Moreover, the performance gap between our method and the baseline widens as the number of experts increases, suggesting enhanced specialization due to finer partitioning of the data distribution.

Notably, despite the absence of capacity constraints on experts during inference, unlike the enforced capacities during training, all experts are actively utilized and receive substantial portions of the data. The consistent improvements across all routed dataset segments confirm that the experts do specialize and they all contribute to the overall improvement instead of only a few dominant experts.




Figure 5: Experts Do specialize. Test perplexity comparison between our method and the dense baseline on the routed dataset segments for the 1.3B parameter model, using mixtures of 4 experts in (a) (trained on 266B tokens), 16 experts in (b) (1T tokens), and 32 experts in (c) (2T tokens). Each bar represents a dataset segment, with the color intensity indicating the percentage of data assigned to that expert – darker shades correspond to a higher proportion of data. Overlapping bars depict the perplexity achieved by the dense baseline (translucent) and our proposed mixture model (opaque). The results demonstrate that all experts specialize effectively on their assigned segments of the data distribution, leading to consistent improvements over the baseline. While the data distribution among experts is not perfectly even – with some experts receiving more data than others – all experts receive a substantial portion of the data. This shows that each expert contributes meaningfully to the overall performance gains.

4 RELATED WORK

Data-parallel training requires synchronization of the gradients after each backward pass, which can be slow if the network bandwidth is limited. In this section, we discuss existing methods that aim to make data parallelism more communication-efficient via reducing the frequency of the synchronization. Since our paper also explores methods for efficient inference without performance degradation, we discuss this branch of research as well.

Distributed optimization To reduce the frequency of synchronisation during training of large models, researchers have explored distributed optimization techniques both theoretically and empirically (McDonald et al., 2009; Stich, 2019; Lian et al., 2018; Wang & Joshi, 2019; McMahan et al., 2017; Lin et al., 2018; Zinkevich et al., 2010). The core principle of these methods is to allow each worker to execute several local training iterations before engaging in global synchronization – each device performs multiple optimizer updates per communication round. This technique also has been studied for LLM pre-training (Diskin et al., 2021; Douillard et al., 2023) and fine-tuning (Borzunov et al., 2022; Liu et al., 2024; Hilmkil et al., 2021; Ro et al., 2022).

While these approaches reduce the frequency of communication, thereby decreasing communication overhead, they often result in degraded model performance compared to every step synchronization (Diskin et al., 2021), or they lack direct performance comparisons in the context of pre-training (Douillard et al., 2023). The degradation in performance is often attributed to the staleness of gradients and the divergence between local and global models due to delayed synchronization (Zhang

et al., 2016; Yu et al., 2019; Mitra et al., 2021). Moreover, to our knowledge, all experiments have been conducted at a small scale ($\leq 400\text{M}$ parameters) (Douillard et al., 2023; Liu et al., 2024; Ro et al., 2022). Therefore, scalable solutions that can reduce communication frequency for LLM training without compromising performance at larger scales are still needed.

Mixture of Experts Mixture of Experts (MoE) models have emerged as a prominent approach to increase model capacity without a proportional increase in inference costs (Jacobs et al., 1991; Jordan & Jacobs, 1994). The key idea behind MoE is to partition the input space into subsets corresponding to specific subtasks, each managed by a specialized expert. By training different experts on distinct tasks, MoE models aim to enhance performance while reducing computational overhead during inference.

Early implementations of MoE employed algorithms such as Gaussian Processes (Tresp, 2000; Theis & Bethge, 2015; Deisenroth & Ng, 2015), Dirichlet Processes (Shahbaba & Neal, 2009), and Support Vector Machines (SVMs) (Collobert et al., 2001) to model these specialized functions. Recent advancements have integrated deep neural networks as experts, providing scalable and flexible implementations suitable for large-scale models.

In natural language processing, researchers have generalized the original MoE idea to enhance large-scale language models. For instance, Shazeer et al. (2017) introduced a Mixture of Experts feed forward layer in which, in contrast to prior works, routing decision is made for every token. This layer takes a token representation as an input and routes it to the best-matched top- k experts. By doing so the number of active parameters is decreased. Similar approaches have been effectively employed in models like the Switch Transformer (Fedus et al., 2022), GShard (Lepikhin et al., 2021), GLaM (Du et al., 2022) and Mixtral of Experts (Jiang et al., 2024) which leverage MoE layers to achieve state-of-the-art results with reduced number of active parameters costs. However, since the routing decision in models like Switch Transformer MoE is made on the token level: (1) all experts must be retained in RAM during both training and inference phases, necessitating substantial memory resources, and (2) a high communication overhead is required for data transfer and gradient synchronization across experts (Lepikhin et al., 2021; Fedus et al., 2022; Riquelme et al., 2021). In that respect this approach is completely orthogonal to our proposed SMALLTALK LM. In fact, each of the experts in our method could be such a token-level mixture model which would result in even smaller inference FLOPs for a specific achieved perplexity albeit with a significantly more complex implementation and higher RAM requirements during inference.

Independent Training of Mixture of Experts Models To train experts fully in parallel, recent research has also explored independent training of the mixture of experts models (Gross et al., 2017; Li et al., 2022; Gururangan et al., 2023). For instance, Gross et al. (2017) proposed to train a gater independently from the experts and use the output of the gater to hard assign the input image to experts. In language processing, Li et al. (2022) and Gururangan et al. (2023) proposed methods that train expert models separately and merge them post-training. However, these approaches rely on full document context or metadata for routing, which can be impractical in real-time applications involving shorter inputs. For instance, Gururangan et al. (2023) focus on performance in text classification tasks, which may not fully reflect real-world scenarios.

5 CONCLUSION

In this work, we present SMALLTALK LM – a method that significantly reduces communication overhead during LLM training. It also outperforms dense model baselines trained with a similar amount of FLOPs and the same data volume and comparable inference costs. In addition, our proposed lightweight routing with short prefixes makes SMALLTALK LM practical for real-world downstream tasks, where we show it largely outperforms the dense baselines.

These findings open up multiple directions for future research, including exploring different routing strategies for data partitioning among experts and investigating further the impact on downstream task performance. Another promising avenue is scaling the mixture to hundreds or even thousands of experts, which could further enhance model performance while benefiting from sparse inference.

REFERENCES

- 540
541
542
543 Alham Fikri Aji and Kenneth Heafield. Sparse communication for distributed gradient descent. In
544 Martha Palmer, Rebecca Hwa, and Sebastian Riedel (eds.), *Proceedings of the 2017 Conference*
545 *on Empirical Methods in Natural Language Processing*, pp. 440–445. Association for Computa-
546 tional Linguistics, September 2017.
- 547
548 Mikel Artetxe, Shruti Bhosale, Naman Goyal, Todor Mihaylov, Myle Ott, Sam Shleifer, Xi Victo-
549 ria Lin, Jingfei Du, Srinivasan Iyer, Ramakanth Pasunuru, Giridharan Anantharaman, Xian Li,
550 Shuohui Chen, Halil Akin, Mandeep Baines, Louis Martin, Xing Zhou, Punit Singh Koura, Brian
551 O’Horo, Jeffrey Wang, Luke Zettlemoyer, Mona Diab, Zornitsa Kozareva, and Veselin Stoyanov.
552 Efficient large scale language modeling with mixtures of experts. In Yoav Goldberg, Zornitsa
553 Kozareva, and Yue Zhang (eds.), *Proceedings of the 2022 Conference on Empirical Methods*
554 *in Natural Language Processing*, pp. 11699–11732. Association for Computational Linguistics,
555 December 2022.
- 556 Alexander Borzunov, Dmitry Baranchuk, Tim Dettmers, Max Ryabinin, Younes Belkada, Artem
557 Chumachenko, Pavel Samygin, and Colin Raffel. Petals: Collaborative inference and fine-tuning
558 of large models. *Association for Computational Linguistics*, 2022.
- 559 James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal
560 Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao
561 Zhang. JAX: composable transformations of Python+NumPy programs, 2018.
- 562
563 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
564 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel
565 Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler,
566 Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray,
567 Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever,
568 and Dario Amodei. Language models are few-shot learners. 33:1877–1901, 2020.
- 569 Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and
570 Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge.
571 *arXiv preprint arXiv:1803.05457*, 2018.
- 572 Ronan Collobert, Samy Bengio, and Yoshua Bengio. A parallel mixture of svms for very large scale
573 problems. *Advances in Neural Information Processing Systems*, 14, 2001.
- 574
575 Common Crawl. Common crawl: Open web data, 2024. Accessed: 2024-09-15.
- 576 Together Computer. Redpajama: An open source recipe to reproduce llama training dataset, 2023.
- 577
578 Marc Deisenroth and Jun Wei Ng. Distributed gaussian processes. In *International conference on*
579 *machine learning*, pp. 1481–1490. PMLR, 2015.
- 580
581 Tim Dettmers. 8-bit approximations for parallelism in deep learning. *arXiv preprint*
582 *arXiv:1511.04561*, 2015.
- 583 Michael Diskin, Alexey Bukhtiyarov, Max Ryabinin, Lucile Saulnier, Anton Sinitsin, Dmitry Popov,
584 Dmitry V Pyrkun, Maxim Kashirin, Alexander Borzunov, Albert Villanova del Moral, et al. Dis-
585 tributed deep learning in open collaborations. *Advances in Neural Information Processing Sys-*
586 *tems*, 34:7879–7897, 2021.
- 587
588 Arthur Douillard, Qixuan Feng, Andrei A Rusu, Rachita Chhaparia, Yani Donchev, Adhiguna
589 Kuncoro, Marc’Aurelio Ranzato, Arthur Szlam, and Jiajun Shen. Diloco: Distributed low-
590 communication training of language models. *arXiv preprint arXiv:2311.08105*, 2023.
- 591
592 Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim
593 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.

- 594 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
595 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.
596 *arXiv preprint arXiv:2407.21783*, 2024.
- 597
598 William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter
599 models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39,
600 2022.
- 601 Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Fos-
602 ter, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muen-
603 nighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang
604 Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. A framework for
605 few-shot language model evaluation, 07 2024.
- 606 Sam Gross, Marc’Aurelio Ranzato, and Arthur Szlam. Hard mixtures of experts for large scale
607 weakly supervised vision. In *Proceedings of the IEEE Conference on Computer Vision and Pat-
608 tern Recognition*, pp. 6865–6873, 2017.
- 609
610 Suchin Gururangan, Margaret Li, Mike Lewis, Weijia Shi, Tim Althoff, Noah A Smith, and Luke
611 Zettlemoyer. Scaling expert language models with unsupervised domain discovery. *arXiv preprint
612 arXiv:2303.14177*, 2023.
- 613 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and
614 Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint
615 arXiv:2009.03300*, 2020.
- 616
617 Tom Henighan, Jared Kaplan, Mor Katz, Mark Chen, Christopher Hesse, Jacob Jackson, Heewoo
618 Jun, Tom B Brown, Prafulla Dhariwal, Scott Gray, et al. Scaling laws for autoregressive generative
619 modeling. *arXiv preprint arXiv:2010.14701*, 2020.
- 620 Agrin Hilmkil, Sebastian Callh, Matteo Barbieri, Leon René Sütfield, Edvin Listo Zec, and Olof
621 Mogren. Scaling federated learning for fine-tuning of large language models. In *International
622 Conference on Applications of Natural Language to Information Systems*, pp. 15–23. Springer,
623 2021.
- 624 Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza
625 Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Train-
626 ing compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- 627
628 Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of
629 local experts. *Neural computation*, 3(1):79–87, 1991.
- 630 Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bam-
631 ford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al.
632 Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.
- 633
634 Michael I Jordan and Robert A Jacobs. Hierarchical mixtures of experts and the em algorithm.
635 *Neural computation*, 6(2):181–214, 1994.
- 636
637 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child,
638 Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language
639 models. *CoRR*, 2020.
- 640 Taku Kudo and John Richardson. SentencePiece: A simple and language independent subword
641 tokenizer and detokenizer for neural text processing. In Eduardo Blanco and Wei Lu (eds.), *Pro-
642 ceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System
643 Demonstrations*, pp. 66–71, Brussels, Belgium, November 2018. Association for Computational
644 Linguistics.
- 645 Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang,
646 Maxim Krikun, Noam Shazeer, and Zhifeng Chen. Gshard: Scaling giant models with conditional
647 computation and automatic sharding. In *International Conference on Learning Representations*,
2021.

- 648 Mike Lewis, Shruti Bhosale, Tim Dettmers, Naman Goyal, and Luke Zettlemoyer. Base layers:
649 Simplifying training of large, sparse models. In *International Conference on Machine Learning*,
650 pp. 6265–6274. PMLR, 2021.
- 651 Margaret Li, Suchin Gururangan, Tim Dettmers, Mike Lewis, Tim Althoff, Noah A Smith, and Luke
652 Zettlemoyer. Branch-train-merge: Embarrassingly parallel training of expert language models.
653 *arXiv preprint arXiv:2208.03306*, 2022.
- 654
- 655 Xiangru Lian, Ce Huang, Yijun Li, and Ji Liu. Asynchronous decentralized parallel stochastic
656 gradient descent. In *Proceedings of the 35th International Conference on Machine Learning*, pp.
657 3043–3052. PMLR, 2018.
- 658
- 659 Yujun Lin, Song Han, Huizi Mao, Yu Wang, and William J Dally. Deep gradient compression: Re-
660 ducing the communication bandwidth for distributed training. *arXiv preprint arXiv:1712.01887*,
661 2017.
- 662
- 663 Yujun Lin, Song Han, Huizi Mao, Yu Wang, and William J Dally. Deep gradient compression:
664 Reducing the communication bandwidth for distributed training. In *International Conference on*
665 *Learning Representations*, 2018.
- 666
- 667 Bo Liu, Rachita Chhaparia, Arthur Douillard, Satyen Kale, Andrei A Rusu, Jiajun Shen, Arthur
668 Szlam, and Marc’Aurelio Ranzato. Asynchronous local-sgd training for language modeling.
arXiv preprint arXiv:2401.09135, 2024.
- 669
- 670 Ryan Mcdonald, Mehryar Mohri, Nathan Silberman, Dan Walker, and Gideon Mann. Efficient
671 large-scale distributed training of conditional maximum entropy models. *Advances in neural*
672 *information processing systems*, 22, 2009.
- 673
- 674 H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas.
675 Communication-efficient learning of deep networks from decentralized data. In *Proceedings of*
676 *the 20th International Conference on Artificial Intelligence and Statistics*, pp. 1273–1282. PMLR,
2017.
- 677
- 678 Praneeth Mitra, Idan Bistriz, Ji Wang, Dan Alistarh, and Wotao Yin. Linear convergence in fed-
679 erated learning: Tackling client heterogeneity and asynchronous updates. In *Advances in Neural*
680 *Information Processing Systems*, volume 34, pp. 14606–14619, 2021.
- 681
- 682 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
683 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-
performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- 684
- 685 Pitch Patarasuk and Xin Yuan. Bandwidth optimal all-reduce algorithms for clusters of workstations.
686 *Journal of Parallel and Distributed Computing*, 69(2):117–124, 2009.
- 687
- 688 Ofir Press, Noah A. Smith, and Mike Lewis. Train short, test long: Attention with linear biases
689 enables input length extrapolation. In *The Tenth International Conference on Learning Represen-*
tations, ICLR 2022, Virtual Event, April 25-29, 2022. OpenReview.net, 2022.
- 690
- 691 Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André
692 Susano Pinto, Daniel Keysers, and Neil Houlsby. Scaling vision with sparse mixture of experts.
693 *Advances in Neural Information Processing Systems*, 34:8583–8595, 2021.
- 694
- 695 Jae Hun Ro, Theresa Breiner, Lara McConnaughey, Mingqing Chen, Ananda Theertha Suresh,
696 Shankar Kumar, and Rajiv Mathews. Scaling language model size in cross-device federated learn-
ing. *arXiv preprint arXiv:2204.09715*, 2022.
- 697
- 698 Babak Shahbaba and Radford Neal. Nonlinear models using dirichlet process mixtures. *Journal of*
699 *Machine Learning Research*, 10(8), 2009.
- 700
- 701 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.
In *International Conference on Learning Representations*, 2017.

- 702 Sebastian U Stich. Local SGD converges fast and communicates little. In *International Conference*
703 *on Learning Representations*, 2019.
704
- 705 Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: En-
706 hanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.
- 707 Rajeev Thakur, Rolf Rabenseifner, and William Gropp. Optimization of collective communication
708 operations in mpich. *The International Journal of High Performance Computing Applications*, 19
709 (1):49–66, 2005.
710
- 711 Lucas Theis and Matthias Bethge. Generative image modeling using spatial lstms. *Advances in*
712 *neural information processing systems*, 28, 2015.
- 713 Volker Tresp. Mixtures of gaussian processes. *Advances in neural information processing systems*,
714 13, 2000.
715
- 716 A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- 717 Jianyu Wang and Gauri Joshi. Adaptive communication strategies to achieve the best error-runtime
718 trade-off in local-update sgd. *Proceedings of Machine Learning and Systems*, 1:212–229, 2019.
719
- 720 Johannes Welbl, Nelson F. Liu, and Matt Gardner. Crowdsourcing multiple choice science ques-
721 tions. In Leon Derczynski, Wei Xu, Alan Ritter, and Tim Baldwin (eds.), *Proceedings of the 3rd*
722 *Workshop on Noisy User-generated Text*, pp. 94–106. Association for Computational Linguistics,
723 September 2017.
- 724 Hao Yu, Sen Yang, and Shenghuo Zhu. Parallel restarted sgd with faster convergence and less
725 communication: Demystifying why model averaging works for deep learning. In *Proceedings of*
726 *the AAAI Conference on Artificial Intelligence*, volume 33, pp. 5693–5700, 2019. doi: 10.1609/
727 aai.v33i01.33015693.
728
- 729 Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a ma-
730 chine really finish your sentence? In *Annual Meeting of the Association for Computational*
731 *Linguistics*, 2019.
- 732 Sixin Zhang, Anna E Choromanska, and Yann LeCun. Deep learning with elastic averaging sgd.
733 *Advances in neural information processing systems*, 28, 2015.
- 734 Wei Zhang, Suyog Gupta, Xiangru Lian, and Ji Liu. Staleness-aware async-sgd for distributed deep
735 learning. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pp.
736 2350–2356, 2016.
737
- 738 Martin Zinkevich, Markus Weimer, Lihong Li, and Alex Smola. Parallelized stochastic gradient
739 descent. *Advances in neural information processing systems*, 23, 2010.
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

756	CONTENTS	
757		
758	1 Introduction	1
759		
760	2 Method	2
761		
762	2.1 Background	2
763		
764	2.2 SMALLTALK LM	3
765		
766	3 Experiments	5
767		
768	3.1 Experimental Setup	6
769		
770	3.2 Language Modeling Results	6
771		
772	3.3 Downstream evaluation	7
773		
774	3.4 Routing analysis	7
775		
776	4 Related Work	9
777		
778	5 Conclusion	10
779		
780	A Language modeling experiments	15
781		
782	A.1 Architecture and training details	15
783		
784	A.2 Notation	15
785		
786	A.3 Computational Cost	17
787		
788	A.3.1 Transformer	17
789		
790	A.3.2 Mixture of LLMs	19
791		
792	A.4 Communication overhead	20
793		
794	B Downstream evaluation	22
795		
796	C Additional Results	23
797		
798	A LANGUAGE MODELING EXPERIMENTS	
799		
800	A.1 ARCHITECTURE AND TRAINING DETAILS	
801		
802	Our architecture is based on the transformer decoder (Vaswani, 2017). We use rotary positional encoding (Su et al., 2024). We use a SentencePiece (Kudo & Richardson, 2018) tokenizer with a vocabulary of 32000 tokens. The model specific parameters for the different sizes are summarized in Table 1. We implemented our EM scheme for training the routers in PyTorch (Paszke et al., 2019). After segmenting the training set, the experts were trained independently using Jax (Bradbury et al., 2018). All training was done in <code>bfloat16</code> . The optimizer states and operations are in <code>float32</code> . The training parameters for specific models are summarised in Table 2.	
803		
804		
805	A.2 NOTATION	
806		
807	For the dense baseline LLM and expert LLM:	
808		
809	• Number of layers (L)	
	• Hidden dimension size (H)	

Model parameters	Role	Hidden size	FFW expansion factor	Layers	Attention heads
335M	expert	1024	4	24	16
1.3B	expert	2048	4	24	16
4.4M	router	96	4	12	12
64M	router	416	4	12	12
110M	router	768	4	12	12

Table 1: Model parameters.

Model	Steps (k)	Tokens (B)	Batch size	# GPUs
335M (dense)	256	133	512	8
335M (4 experts)	256	133	128	8
335M (dense)	512	266	512	32
335M (8 experts)	256	266	128	8
335M (dense)	1024	532	512	32
335M (16 experts)	256	532	128	8
1.3B (dense)	512	266	512	32
1.3B (4 experts)	512	266	128	8
1.3B (dense)	1024	1000	1024	64
1.3B (16 experts)	512	1000	128	8
1.3B (dense)	1024	2000	2048	128
1.3B (32 experts)	512	2000	128	8
4.4M (router)	128	4	32	1

Table 2: Training parameters.

- Number of attention heads (A)
- Sequence length (S)
- Batch size (B)
- Vocabulary size (V)
- Feedforward network dimension (D_{ff})

For the router LLM:

- Prefix length for routing (M)
- Number of layers (L_r)
- Hidden dimension size (H_r)
- Number of attention heads (A_r)
- Batch size (B_r)
- Feedforward network dimension (D_{rff}) – parameters specific to the router model
- Number of training tokens per router between communications T

We also define N_{steps} , N_{steps_expert} , and N_{steps_router} as the number of training steps for the dense baseline, the expert in the mixture, and the router, respectively. The sequence length (S) and vocabulary size (V) are consistent across the dense baseline, router, and expert models. Finally, we denote E as the number of experts (routers).

864 A.3 COMPUTATIONAL COST

865 A.3.1 TRANSFORMER

866 To estimate the computational cost of Transformer model during training and inference, we calcu-
867 lated the total floating-point operations (FLOPs) based on the model’s architectural parameters.

868 For **training**, we computed the FLOPs per training step by summing the contributions from the
869 embedding layer, multi-head attention (MHA), feedforward layers (FFN), and the output projection
870 layer. The calculations are as follows:

- 871 • **Embedding Layer:** Although embeddings are typically implemented as lookups with neg-
872 ligible computational cost, for completeness, we estimate the FLOPs as:

$$873 \text{FLOPs}_{\text{emb}} = B \times S \times H$$

- 874 • **Multi-Head Attention (MHA):**

- 875 1. Linear Projections (Queries, Keys, Values):

$$876 \text{FLOPs}_{\text{proj}} = 3 \times 2 \times B \times S \times H \times H = 6 \times B \times S \times H^2$$

- 877 2. Scaled Dot-Product Attention:

$$878 \text{FLOPs}_{\text{attn}} = \text{FLOPs}_{\text{QK}} + \text{FLOPs}_{\text{V}} = 2 \times B \times S^2 \times H + 2 \times B \times S^2 \times H = 4 \times B \times S^2 \times H$$

- 879 3. Output Projection:

$$880 \text{FLOPs}_{\text{out.proj}} = 2 \times B \times S \times H \times H$$

881 Total Multi-Head Attention (MHA):

$$\begin{aligned} 882 \text{FLOPs}_{\text{MHA}} &= \text{FLOPs}_{\text{proj}} + \text{FLOPs}_{\text{attn}} + \text{FLOPs}_{\text{out.proj}} \\ 883 &= 6 \times B \times S \times H^2 + 4 \times B \times S^2 \times H + 2 \times B \times S \times H^2 \\ 884 &= 8 \times B \times S \times H^2 + 4 \times B \times S^2 \times H \end{aligned}$$

- 885 • **Feedforward Network (FFN):**

$$886 \text{FLOPs}_{\text{FFN}} = 2 \times 2 \times B \times S \times H \times D_{\text{ff}} = 4 \times B \times S \times H \times D_{\text{ff}}$$

- 887 • **Output projection:**

$$888 \text{FLOPs}_{\text{out}} = \text{FLOPs}_{\text{out.proj}} + \text{FLOPs}_{\text{softmax}} = 2 \times B \times S \times H \times V + 3 \times B \times S \times V$$

889 **Total FLOPs per Layer:**

$$890 \text{FLOPs}_{\text{layer}} = (\text{FLOPs}_{\text{MHA}} + \text{FLOPs}_{\text{FFN}}) \times L$$

891 **Total Forward Pass FLOPs per Step:**

$$892 \text{FLOPs}_{\text{forward}} = \text{FLOPs}_{\text{emb}} + \text{FLOPs}_{\text{layer}} + \text{FLOPs}_{\text{out}}$$

893 **Total Training FLOPs per Step:**

894 The backward pass is approximately twice as computationally intensive as the forward pass. There-
895 fore:

$$896 \text{FLOPs}_{\text{train.step}} = \text{FLOPs}_{\text{forward}} + 2 \times \text{FLOPs}_{\text{forward}} = 3 \times \text{FLOPs}_{\text{forward}}$$

897 **Total Training FLOPs:**

898 Multiplying the FLOPs per training step by the total number of training steps:

$$899 \text{Total Training FLOPs} = \text{FLOPs}_{\text{train.step}} \times N_{\text{steps}} = 3 \times \text{FLOPs}_{\text{forward}} \times N_{\text{steps}}$$

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

Total Training FLOPs:

$$\begin{aligned} \text{Total Training FLOPs} = 3 \times N_{\text{steps}} \times & \left[B \times S \times H \right. \\ & + L \times \left(8 \times B \times S \times H^2 + 4 \times B \times S^2 \times H + 4 \times B \times S \times H \times D_{\text{ff}} \right) \\ & \left. + 2 \times B \times S \times H \times V + 3 \times B \times S \times V \right] \quad (10) \end{aligned}$$

Total Inference FLOPs:

$$\begin{aligned} \text{Total Inference FLOPs} = B \times & \left[S \times H \right. \\ & + L \times \left(8 \times S \times H^2 + 4 \times S^2 \times H + 4 \times S \times H \times D_{\text{ff}} \right) \\ & \left. + 2 \times S \times H \times V + 3 \times S \times V \right] \quad (11) \end{aligned}$$

972 A.3.2 MIXTURE OF LLMs

973 The total training FLOPs for the mixture model consist of four main components: training routers,
974 shading data for routers, training experts, and shading data for the experts.

975 For a model trained with batch size B over N_{steps} steps, the total number of sequences shaded can
976 be estimated as $N_{\text{steps}} \times B \times E$, where each model receives approximately $N_{\text{steps}} \times B$ sequences
977 during training.

978 Thus, the FLOPs required for shading can be calculated as the inference FLOPs with an effective
979 batch size of $N_{\text{steps}} \times B \times E$ (see Equation 11).

980 Therefore **Total Training FLOPs:**

981 Total Training FLOPs = Training FLOPs Routers + Training FLOPs Experts = (12)

$$\begin{aligned}
 &+ 3 \times N_{\text{steps_router}} \times \left[B_r \times S \times H_r \right. \\
 &+ L_r \times (8 \times B_r \times S \times H_r^2 + 4 \times B_r \times S^2 \times H_r + 4 \times B_r \times S \times H_r \times D_{\text{rff}}) \\
 &\left. + 2 \times B_r \times S \times H_r \times V + 3 \times B_r \times S \times V \right] \times E \quad (\text{Eq. 13: Training Routers})
 \end{aligned}$$

982 (13)

$$\begin{aligned}
 &+ \left(N_{\text{steps_router}} \times B_r \times E \right) \times \left[M \times H_r \right. \\
 &+ L_r \times (8 \times M \times H_r^2 + 4 \times M^2 \times H_r + 4 \times M \times H_r \times D_{\text{rff}}) \\
 &\left. + 2 \times M \times H_r \times V + 3 \times M \times V \right] \times E \quad (\text{Eq. 14: Shading data for routers})
 \end{aligned}$$

983 (14)

$$\begin{aligned}
 &3 \times N_{\text{steps_expert}} \times \left[B \times S \times H \right. \\
 &+ L \times (8 \times B \times S \times H^2 + 4 \times B \times S^2 \times H + 4 \times B \times S \times H \times D_{\text{ff}}) \\
 &\left. + 2 \times B \times S \times H \times V + 3 \times B \times S \times V \right] \times E \quad (\text{Eq. 15: Training Experts})
 \end{aligned}$$

984 (15)

$$\begin{aligned}
 &+ \left(N_{\text{steps_expert}} \times B \times E \right) \times \left[M \times H_r \right. \\
 &+ L_r \times (8 \times M \times H_r^2 + 4 \times M^2 \times H_r + 4 \times M \times H_r \times D_{\text{rff}}) \\
 &\left. + 2 \times M \times H_r \times V + 3 \times M \times V \right] \times E \quad (\text{Eq. 16: Shading data for experts})
 \end{aligned}$$

985 (16)

986 In the above formula:

- 987 • Equations 15 and 13 correspond to the total training FLOPs for E experts and E routers
988 respectively.
- 989 • Equations 16 and 14 represents the inference cost of E routers on B batches with a prefix
990 length M used to partition the data for experts and routers respectively.

991 Table 3 presents the training and inference costs for our proposed mixture of language models com-
992 pared to the dense baseline LLMs.

The expert balancing during training introduces only a negligible computational overhead relative to both training and inference costs. Therefore, it is excluded from our FLOPs calculations.

Model	Perplexity	Training cost	Inference cost
335M (dense)	11.78	31.02	0.79
335M (4 experts)	10.78 ↓ 8.49%	31.02 + 0.22 ↑ 0.07%	0.79 + 0.01 ↑ 1.27%
335M (dense)	11.25	62.03	0.79
335M (8 experts)	10.20 ↓ 9.33%	62.03 + 0.75 ↑ 1.20%	0.79 + 0.02 ↑ 2.53%
335M (dense)	10.80	124.06	0.79
335M (16 experts)	9.64 ↓ 10.74%	124.06 + 2.71 ↑ 2.19%	0.79 + 0.04 ↑ 5.06%
335M (dense)	10.5	248.12	0.79
335M (32 experts)	9.07 ↓ 13.62%	248.12 + 10.28 ↑ 4.14%	0.79 + 0.08 ↑ 10.13%
1.3B (dense)	9.10	221.33	2.81
1.3B (4 experts)	8.75 ↓ 3.85%	221.33 + 0.36 ↑ 0.16%	2.81 + 0.01 ↑ 0.36%
1.3B (dense)	8.48	885.32	2.81
1.3B (16 experts)	7.42 ↓ 12.53%	885.32 + 4.87 ↑ 0.55%	2.81 + 0.04 ↑ 1.42%
1.3B (dense)	8.20	1770.65	2.81
1.3B (32 experts)	6.76 ↓ 17.56%	1770.65 + 18.94 ↑ 1.07%	2.81 + 0.08 ↑ 2.85%

Table 3: **Performance gain versus computational overhead.** The table compares performance improvements and computational costs for models with 335M (trained with 4, 8, 16, and 32 experts) and 1.3B parameters (trained with 4, 16, and 32) experts. Training cost is measured in 10^{19} FLOPs, inference cost is measured in 10^{12} FLOPs. The dense baselines and corresponding mixture models were trained on the same data volume.

Regarding Computational Overhead As demonstrated in Table 3, the 1.3B parameter model with 32 experts achieves significant performance improvements almost cost-free, gaining nearly 18% in perplexity compared to the dense baseline while incurring only 1% additional cost during training and less than 3% during inference. Furthermore, a 335M parameter mixture of 32 experts reaches the same performance level as the 1.3B parameter model, but with a comparable training budget and **three times lower inference cost**. Despite the 335M parameter model incurring a higher percentage cost during inference, this computational overhead is not the definitive minimum necessary for effective routing: as detailed in § 3.4, our results indicate that the routing overhead can be significantly reduced by utilizing smaller routers and shorter prefix sizes. This suggests that further reductions in computational costs, without compromising routing quality, are feasible and present a valuable direction for future research.

A.4 COMMUNICATION OVERHEAD

In this section, we provide an estimate of the communication overhead involved in our model training compared to a dense baseline trained with standard distributed data parallelism. We quantify both the volume of data transmitted and the frequency of communication required during training. For simplicity of computation, we assume that all collective communications are bandwidth-optimal, meaning that the total number of bytes transferred per node is approximately $2K$, where K is the size of the message (Thakur et al., 2005; Patarasuk & Yuan, 2009).

Mixture of LLMs Every router performs all gather operation to share the loss on the dataset chunk approximately every $T = 45\text{M}$ training tokens, as outlined in 3.2. This frequency was primarily determined by implementation specifics related to data storage.

We can define the number of tokens processed per step by a router as $N_{\text{tokens_per_step}}$, which can be calculated as:

$$N_{\text{tokens_per_step}} \leq S \times B_r$$

This leads to the minimum number of steps between communications, $N_{\text{steps.per.com}}$, given by:

$$N_{\text{steps.per.com}} \geq \frac{T}{S \times B_r}$$

The maximum number of communication events, N_{comm} , during the router’s training can then be estimated as:

$$N_{\text{comm}} \leq \frac{N_{\text{steps.router}}}{N_{\text{steps.per.comm}}} = \frac{N_{\text{steps.router}} \times S \times B_r}{T}$$

As shown in App. Table 2 we train routers for 128,000 steps with batch size 32. We use context size 1024. Therefore, $N_{\text{comm}} \approx 94 < 1 \times 10^2$.

The total amount of data sent and received by each node (router) can be calculated as the total number of sequences multiplied by the number of bytes per sequence. We consider that the loss is transferred in `float16`, which requires 2 bytes per element. The total number of sequences can be estimated as $\frac{T \times N}{S}$.

Each router needs to send and receive the loss values for each sequence to and from other routers, resulting in a factor of 2 for the data sent and received. Therefore, the data per router can be approximated as:

$$\text{Data per router} \approx 2 \times 2 \times \frac{T \times E}{S}$$

Given $T = 45 \times 10^6$ – training tokens per router between communication steps, $E \leq 32$ – number of experts and $S = 1,024$ – sequence length, we can estimate the data transferred per node as:

$$\text{Data per router} \leq 2 \times 2 \times \frac{T \times E}{S} = 4 \times \frac{45 \times 10^6 \times 32}{1,024} = 5.625\text{MB}$$

Once routers are fully trained, they communicate only when necessary, such as when experts require new data for training.

For instance, consider that the message size is K bytes, equivalent to loss array of $K/2$ sequences in `float16` precision. The frequency of communication after routers are fully trained is determined by the batch size of the expert B and the number of experts E . Specifically, the number of steps each expert makes between communications can be estimated as:

$$N_{\text{steps.per.comm}} \approx \frac{K}{2 \times B \times E}$$

Thus, communication can be scheduled approximately every $\frac{K}{4 \times B \times E}$ expert’s training steps. Assuming $K = 2^{31} - 1$ bytes, $B = 128$ as used in the majority of our experiments, and $E \in \{4, 8, 16, 32\}$ (see Appendix 2), the communication interval becomes:

$$\frac{2^{31} - 1}{2 \times 128 \times E} \approx \frac{2^{23}}{E} \leq 2^{18} \quad \forall E \in \{4, 8, 16, 32\} \quad (17)$$

Comparison with Distributed Training Let us assume that we train a model with a batch size of B on a single host, where W represents the model size in parameters. For simplicity, we consider only data parallelism and assume that the model weights can fit on one GPU. To compare communication costs, we examine the additional expense incurred when increasing throughput by a factor of E , i.e., utilizing E hosts instead of one.

In regular distributed training, scaling to E hosts necessitates copying all weights to each host and performing gradient synchronization at every training step. Assuming each parameter is represented by 32 bits (as optimizers typically operate in `float32`), the amount of data transferred to and from each node during gradient synchronization can be calculated using the bandwidth-optimal algorithm:

$$\text{Data per node} = 2 \times W \times 4$$

Here, $W \times 4$ represent the size of the model gradients in bytes.

If $W = 1.3 \times 10^9$ parameters, then the amount of data transferred per node per training step is:

$$\text{Data per node} = 2 \times 1.3 \times 10^9 \times 4 = 10.4\text{GB}$$

Therefore, in regular distributed training, each node must send and receive **gigabytes** of data **at every step**. In contrast, our approach offers flexible communication strategies. We can transfer large volumes of data only a few times during training (see Equation 17), or optimise for more frequent communications with smaller data sizes. This flexibility allows us to adapt based on network latency and other infrastructure constraints.

B DOWNSTREAM EVALUATION

For the evaluation, we utilized the *lm-eval-harness* software (Gao et al., 2024). For the computation of perplexity, we compute the perplexity using the format: “Question: {question}. Answer: {answer}.”

Task	32 experts, 2T	dense, 2T
mmlu abstract algebra	0.1400	0.2300
mmlu astronomy	0.3553	0.2829
mmlu business ethics	0.4900	0.5300
mmlu college computer science	0.2500	0.2900
mmlu college mathematics	0.1500	0.1400
mmlu computer security	0.3500	0.3800
mmlu elementary mathematics	0.2672	0.2487
mmlu formal logic	0.3016	0.2937
mmlu global facts	0.3100	0.2800
mmlu high school geography	0.3586	0.3586
mmlu high school macroeconomics	0.3128	0.2872
mmlu high school microeconomics	0.3193	0.3193
mmlu high school statistics	0.2917	0.2593
mmlu human aging	0.4081	0.3857
mmlu human sexuality	0.3969	0.3893
mmlu international law	0.2314	0.2314
mmlu logical fallacies	0.3067	0.2945
mmlu machine learning	0.2589	0.2679
mmlu medical genetics	0.3900	0.3500
mmlu moral disputes	0.2803	0.2890
mmlu moral scenarios	0.2380	0.2380
mmlu nutrition	0.3039	0.2516
mmlu prehistory	0.4167	0.3858
mmlu professional accounting	0.2872	0.2518
mmlu professional medicine	0.3566	0.2904
mmlu professional psychology	0.3203	0.3186
mmlu world religions	0.5731	0.4503
mmlu anatomy	0.4667	0.3926
mmlu clinical knowledge	0.3434	0.2755
mmlu college biology	0.4028	0.3819
mmlu college chemistry	0.2800	0.2600
mmlu college medicine	0.2832	0.2659
mmlu college physics	0.1863	0.1961
mmlu conceptual physics	0.3787	0.3787

Table 4: **Downstream Evaluation Results (Part 1)**. This table presents a comparison in accuracy on downstream tasks between a 1.3B-parameter, 32-expert model and a 1.3B-parameter dense model. Both models were trained using the 2T tokens and nearly the same training FLOPs, with only a 1% difference. The inference cost of the expert model increased by less than 3% compared to the dense baseline FLOPs.

Task	32 experts, 2T	dense, 2T
mmlu econometrics	0.2719	0.2018
mmlu electrical engineering	0.2552	0.2621
mmlu high school biology	0.3290	0.3097
mmlu high school chemistry	0.2167	0.1970
mmlu high school computer science	0.2400	0.2400
mmlu high school european history	0.2848	0.3212
mmlu high school government and politics	0.4093	0.3834
mmlu high school mathematics	0.1519	0.1444
mmlu high school physics	0.2980	0.2715
mmlu high school psychology	0.4642	0.4330
mmlu high school us history	0.3382	0.3137
mmlu high school world history	0.2827	0.2996
mmlu jurisprudence	0.2407	0.2222
mmlu management	0.3786	0.3981
mmlu marketing	0.4915	0.4915
mmlu miscellaneous	0.5121	0.4994
mmlu philosophy	0.2958	0.3119
mmlu professional law	0.2634	0.2536
mmlu public relations	0.3545	0.4273
mmlu security studies	0.3224	0.3020
mmlu sociology	0.2786	0.2786
mmlu us foreign policy	0.3300	0.3400
arc challenge	0.3200	0.2892
arc easy	0.6780	0.6549
hellaswag	0.4912	0.4663
sciq	0.8950	0.9130

Table 5: Downstream Evaluation Results (Part 2)

C ADDITIONAL RESULTS

Effect of Prefix Length on Routing Performance We also conducted experiments by partitioning the data using a shorter prefix length of $M = 32$ tokens during training, whereas in our main results (see § 3), we used a prefix length of $M = 256$ tokens.

We observed that training with a smaller prefix length resulted in better expert performance when routing during inference was done using a small number of tokens (see Figure 6). This suggests that using a shorter prefix during training enhances the router’s ability to make effective decisions with limited context.

We did not take any action to make sure that the router models could extrapolate to longer sequences than seen in training (e.g. interpolating positional encodings, using ALiBi (Press et al., 2022), etc.). As expected, we observe a performance degradation when making routing decisions for sequences longer than the ones used for training.




Figure 6: Test perplexity as a function of routing prefix length during inference for 1.3B parameter models with 16 experts. The orange curve represents a prefix length of 32, while the blue curve corresponds to a prefix length of 256.