

# $\infty$ -MoE: Generalizing Mixture of Experts to Infinite Experts

Shota Takashiro, Takeshi Kojima,

Shohei Taniguchi, Yusuke Iwasawa, Yutaka Matsuo

The University of Tokyo, Hongo 7-3-1, Bunkyo-ku, Tokyo, 113-8656 Japan

{takashiro, taniguchi, t.kojima, iwasawa, matsuo}@weblab.t.u-tokyo.ac.jp

## Abstract

The Mixture of Experts (MoE) selects a few feed-forward networks (FFNs) per token, achieving an effective trade-off between computational cost and performance. In conventional MoE, each expert is treated as entirely independent, and experts are combined in a discrete space. As a result, when the number of experts increases, it becomes difficult to train each expert effectively. To stabilize training while increasing the number of experts, we propose  $\infty$ -MoE that selects a portion of the parameters of large FFNs based on continuous values sampled for each token. By considering experts in a continuous space, this approach allows for an infinite number of experts while maintaining computational efficiency. Experiments show that a GPT-2 Small-based  $\infty$ -MoE model, with 129M active and 186M total parameters, achieves comparable performance to a dense GPT-2 Medium with 350M parameters. Adjusting the number of sampled experts at inference time allows for a flexible trade-off between accuracy and speed, with an improvement of up to 2.5% in accuracy over conventional MoE.

## 1 Introduction

Large language models (LLMs) have recently achieved remarkable performance across a broad range of natural-language processing (NLP) tasks, such as machine translation, question answering, and code generation (Chen et al., 2021; Liu et al., 2021). These advances are primarily driven by scaling up model parameters, training data, and compute resources (Kaplan et al., 2020). However, simply increasing model size leads to substantial computational and memory overheads, motivating research into more efficient strategies for scaling.

Mixture of Experts (MoE) (Shazeer et al., 2017) stands out for its ability to expand parameter count while maintaining relatively low per-token compute costs. By routing each input to a subset

of specialized experts, MoE-based architectures can efficiently store large amounts of knowledge sparsely (Dai et al., 2024; Jiang et al., 2024). Recent large-scale models such as DeepSeek (Dai et al., 2024), Mistral (Jiang et al., 2024), and Phi (Abdin et al., 2024) have successfully adopted MoE designs, demonstrating that sparse routing can significantly improve performance without incurring prohibitive computational expense.

A notable trend in recent MoE research is to aggressively increase the number of experts for finer-grained specialization. Empirical evidence shows that larger expert pools improve overall capacity and often yield higher accuracy with similar or reduced compute costs (Fedus et al., 2022; Lepikhin et al., 2020). For instance, PEER (He, 2024) can handle millions of experts, and recent theoretical work (Clark et al., 2022) confirms that MoE performance scales predictably with the expert count.

Following this trend, a natural question arises: **can we achieve even better performance by further increasing the number of experts to infinity?** In principle, having more experts should allow for even more specialized representations, potentially boosting generalization across diverse tasks.

We introduce  $\infty$ -MoE, which moves from a discrete set of experts to a continuous domain, allowing theoretically unbounded expert capacity. In this framework, each input samples from a continuum of experts, taking the concept of “increasing experts” to the extreme. Despite the potential for an infinite number of experts, our proposed  $\infty$ -MoE remains computationally tractable because of its sparse activation of only a small number of sampled experts at any given time. This design preserves the efficiency of sparse routing while offering significantly enhanced model capacity. Through experiments on GPT-2 Small and Medium (Radford et al., 2019), we observed that the GPT-2 Small-based  $\infty$ -MoE variant (129 million active parameters, 186 million total) achieves performance comparable to

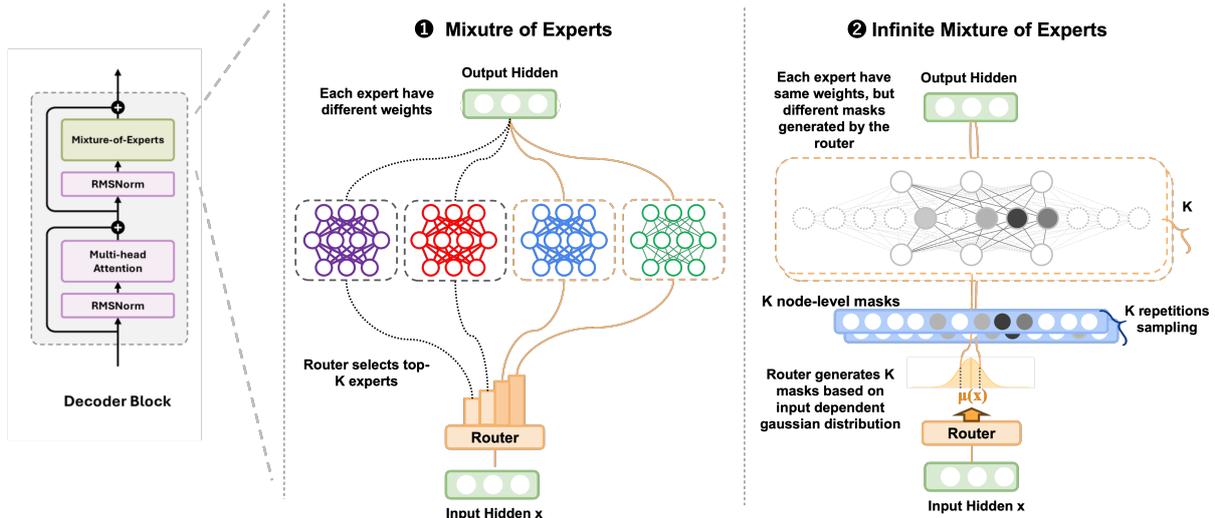


Figure 1: Overview of the proposed Infinite Mixture of Experts ( $\infty$ -MoE). The router outputs a continuous distribution over the expert space, and each sample selects a unique expert.

that of a dense GPT-2 Medium model with 350 million parameters. Furthermore, increasing the number of samples during inference yields additional accuracy gains, and reducing it still maintains a 2.5% accuracy improvement over standard MoE, enabling flexible tradeoffs between speed and accuracy.

## 2 Related Work

MoE was first proposed to split a problem space into multiple specialized expert networks (Jacobs et al., 1991), and has lately gained popularity for LLMs. A central advantage in LLMs is that routing each token to just a few experts can greatly expand parameter capacity without a matching increase in compute (Shazeer et al., 2017; Lepikhin et al., 2020; Fedus et al., 2022). For instance, GShard (Lepikhin et al., 2020) and Switch Transformer (Fedus et al., 2022) employ sparse expert activation to train models with hundreds of billions of parameters, though they typically rely on a small expert pool (16 to a few hundred) that restricts specialization.

Recent work addresses this issue by substantially raising the expert count. PEER (He, 2024) scales up to a million experts, demonstrating richer specialization by novel routing mechanisms. Theoretically, increasing experts improves performance without linearly increasing compute (Clark et al., 2022; Ludziejewski et al., 2024), but router overhead can grow large or over-compressed experts may degrade accuracy (Ludziejewski et al., 2024).

## 3 Proposed Method

This section presents our  $\infty$ -MoE framework. We first introduce a generalized MoE formulation for the standard case, then detail the  $\infty$ -MoE model, which extends MoE to a continuous expert space.

### 3.1 Mixture of Experts

Let  $\mathcal{Z} = \{1, 2, \dots, n\}$  be a discrete index set of  $n$  experts. Let  $x \in \mathbb{R}^{d_{in}}$  denote the input. Each expert is a function:

$$f(x, i) : \mathbb{R}^{d_{in}} \times \mathcal{Z} \rightarrow \mathbb{R}^{d_{out}},$$

where  $i \in \mathcal{Z}$  indexes the expert. A router produces a probability distribution  $p(i|x)$  over experts.

The MoE output is the expected expert output:

$$y = \sum_{i=1}^n p(i|x) f(x, i) \quad (1)$$

**Connection to Standard MoE.** Standard MoE can be seen as a special case where the general expert function  $f(x, i)$  simply selects the  $i$ -th expert from a set of  $n$  pre-defined expert functions,  $\{e_1(x), \dots, e_n(x)\}$ ; that is,  $f(x, i) = e_i(x)$ . The router typically uses a softmax function to compute the probability of selecting expert  $i$ :

$$p(i|x) = \text{softmax}(\text{Top}K(g(x)))_i \quad (2)$$

where  $g(x) \in \mathbb{R}^n$  is a vector of scores produced by the router network. With a top- $k$  operation selecting a subset  $K$  of experts, the final output is

$$y = \sum_{i \in K} p(i|x) e_i(x). \quad (3)$$

Table 1: Zero-shot performance on various benchmarks (BoolQ (Clark et al., 2019), HellaSwag (Zellers et al., 2019), WinoGrande (Sakaguchi et al., 2021), ARC-e/c (Boratto et al., 2018), OpenBookQA (Banerjee et al., 2019), and RACE-high (Lai et al., 2017)). “Active/Total Param” indicates the approximate number of parameters used during forward vs. total parameters.

Model	Active/Total Param	BoolQ(↑)	HellaSwag(↑)	WinoGrande(↑)	ARC-e(↑)	ARC-c(↑)	OBQA(↑)	RACE-high(↑)	Avg(↑)
<b>GPT-2 Small</b>									
Dense	124M/124M	60.1	29.2	50.8	43.1	<b>19.4</b>	15.2	51.3	38.5
Switch Transformer	124M/181M	60.1	29.2	51.2	43.1	18.0	14.4	51.3	38.2
MoE	124M/181M	<b>60.5</b>	29.5	51.5	44.6	18.5	15.8	51.3	38.8
∞-MoE	129M/186M	59.6	<b>29.8</b>	<b>54.2</b>	<b>46.0</b>	18.9	<b>17.6</b>	<b>52.3</b>	<b>39.8</b>
<b>GPT-2 Medium</b>									
Dense	350M/350M	60.7	31.4	48.8	47.1	20.1	17.6	53.1	39.8
Switch Transformer	350M/556M	58.4	31.5	50.0	48.0	20.0	16.2	55.2	39.9
MoE	350M/556M	<b>59.3</b>	32.7	50.7	48.3	20.6	17.8	52.7	40.3
∞-MoE	362M/568M	56.6	<b>33.7</b>	<b>51.6</b>	<b>49.7</b>	<b>21.5</b>	<b>18.8</b>	<b>57.0</b>	<b>41.3</b>

This result clearly demonstrates that the standard MoE is a special case of this discrete formulation.

### 3.2 ∞-MoE: Infinite Experts

∞-MoE extends the discrete MoE to a continuous, potentially uncountably infinite, expert space  $\mathcal{Z} \subseteq \mathbb{R}^{d_z}$ . The router defines a probability density  $p(z|x)$  over  $\mathcal{Z}$ . The model output is

$$y = \int_{\mathcal{Z}} p(z|x) f(x, z) dz \quad (4)$$

We approximate this integral by Monte Carlo sampling: We sample  $z \sim p(z|x)$  and use  $f(x, z)$  as an estimator of  $y$ .

**Router Design.** We use a Gaussian density for the router:

$$p(z|x) = \mathcal{N}(z|\mu(x), \Sigma(x)), \quad (5)$$

where a small neural network predicts  $\mu(x)$  and  $\Sigma(x)$  (i.e., all off-diagonal entries are zero) from  $x$ . During training, we sample  $z^{(k)} \sim p(z|x)$   $K$  times ( $k = 1, \dots, K$ ), allowing the router to learn to allocate probability mass to appropriate regions of  $\mathcal{Z}$ .

**Expert Design.** We treat  $z$  as a continuous expert index sampled from the router. Intuitively, each distinct value of  $z$  corresponds to a different expert in an infinite expert space. Our feed-forward network (FFN) is then modulated by a mask that “turns off” certain neurons in the intermediate layer, allowing the model to dynamically select which subset of parameters is active.

Formally, let  $W_z \in \mathbb{R}^{d_{\text{ff}} \times d_z}$ . Given  $z$  sampled from Equation 5, we apply a top- $N\%$  operator on intermediate neurons  $\hat{m}_i = W_z z$ , which keeps

only the largest  $N\%$  of nodes and sets the rest to 0:

$$\text{mask}(z)_i = \begin{cases} \hat{m}_i & \text{if } \hat{m}_i \text{ is top } N\% \text{ values,} \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

Because the retained entries preserve their original values, the resulting mask is partially “soft” for the selected positions, while all other positions become strictly zero. Given this mask, the expert output  $f(x, z)$  is computed as

$$f(x, z) = W_2 \left( \text{Act}(W_1 x) \odot \text{mask}(z) \right), \quad (7)$$

where  $\text{Act}(\cdot)$  is a nonlinear activation,  $\odot$  is element-wise multiplication, and  $W_1 \in \mathbb{R}^{d_{\text{ff}} \times d_{\text{in}}}$ ,  $W_2 \in \mathbb{R}^{d_{\text{out}} \times d_{\text{ff}}}$  are learnable weight matrices. Through this mechanism, each sampled  $z$  effectively activates a distinct subset of the FFN’s neurons, mirroring the sparsity in conventional MoE models but generalized to infinite experts.

## 4 Experiments

We evaluated the effectiveness of ∞-MoE using GPT-2 Small (~124 million parameters) and GPT-2 Medium (~350 million parameters) on a broad range of natural-language understanding tasks.

### 4.1 Setup

**Data.** We pre-trained our models on a large-scale web corpus called FineWeb (Penedo et al., 2024), from which we extracted 10 billion tokens. For fine-tuning or direct evaluation, we used the zero-shot setting on standard NLP benchmarks.

**Compared Methods.** We compared four architectures:

- **Dense (FFN):** A standard transformer with a single FFN layer shared by all inputs.

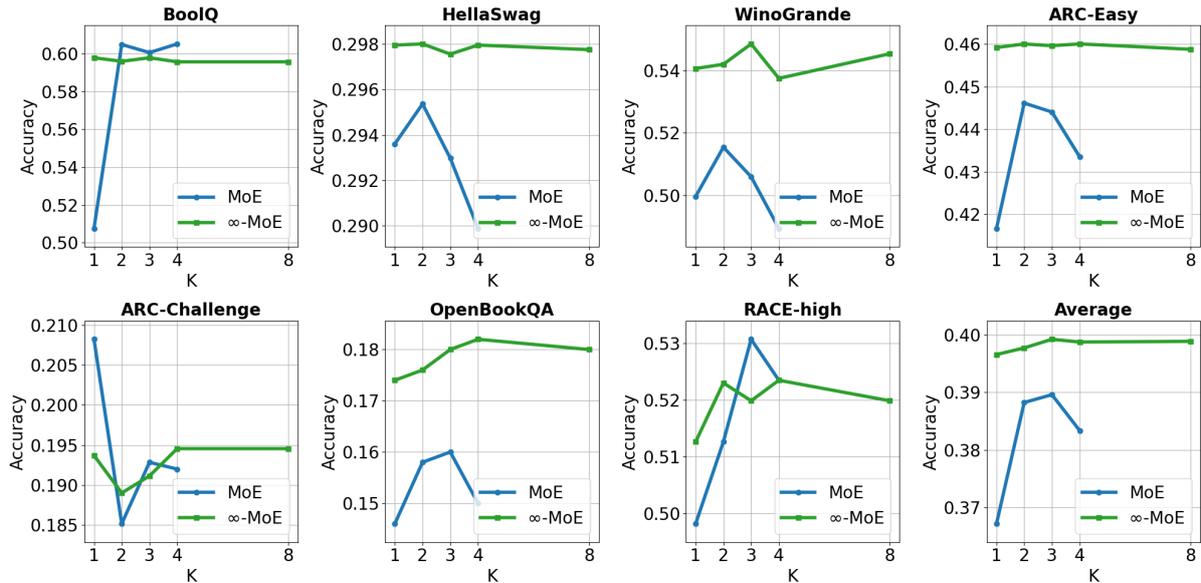


Figure 2: Comparison of MoE and  $\infty$ -MoE models on several tasks while varying the number of experts  $K \in \{1, 2, 3, 4, 8\}$ . For GPT-2 Small,  $K = 2$  yields 124 million active parameters.  $\infty$ -MoE consistently achieves strong accuracy across a wide range of  $K$ , even with fewer experts.

- **Switch Transformer (Top-1):** Routes each token to exactly one expert.
- **MoE (Top-2):** A classic sparse MoE setting that activates the top-2 experts for each token. In this configuration, the total number of experts is fixed at 4.
- **$\infty$ -MoE:** Our proposed method with an infinite expert space. During both training and testing, two samples are drawn (i.e.,  $K = 2$ ); with one sample, only 25% of the overall expert space is active.

## 4.2 Main Results

Table 1 presents the zero-shot performance of each model on GPT-2 Small and GPT-2 Medium. Across all tasks,  $\infty$ -MoE consistently outperformed the Dense baseline, Switch Transformer, and standard MoE. Notably, for GPT-2 Small,  $\infty$ -MoE achieved the highest average score of 39.8 versus 38.5 (Dense), 38.2 (Switch), and 38.8 (MoE). We observed similar improvements with the GPT-2 Medium variant, where  $\infty$ -MoE again attained the best average accuracy (41.3).

## 5 Ablations

### 5.1 Scaling with sampling ( $K$ )

Figure 2 compares  $\infty$ -MoE with standard MoE across multiple tasks by varying  $K$ . In the conventional setup, increasing  $K$  can improve accuracy but may also introduce instability at high values.

By contrast,  $\infty$ -MoE scales more smoothly with  $K$ , yielding robust gains and maintaining strong performance even at lower  $K$  (achieving a 2.5% improvement over standard MoE). Moreover, treating experts as a continuous space enables flexible inference, allowing users to adjust  $K$  according to hardware constraints or latency requirements.

These results demonstrate that  $\infty$ -MoE combines the expressiveness of an unbounded expert ensemble with the efficiency of sparse MoE, making it well suited to a variety of runtime conditions.

### 5.2 Scaling with Dataset Size

To evaluate the effectiveness of our proposed  $\infty$ -MoE method under increasing dataset sizes, we conducted experiments using a GPT-2 Small architecture as the base model. We measured the accuracy on the HellaSwag dataset, progressively increasing the training data size in increments of 10 billion tokens up to 100 billion. The results were plotted as shown in Figure 3.

### 5.3 Dimension of Continuous Index $z$

We varied the dimensionality of the continuous expert index  $z$  and report the perplexity in Table 2. Increasing the dimension improved performance up to  $d_z = 128$ , but larger dimensions provided no further gain. This result suggests that a moderate index dimension is sufficient to represent diverse experts.

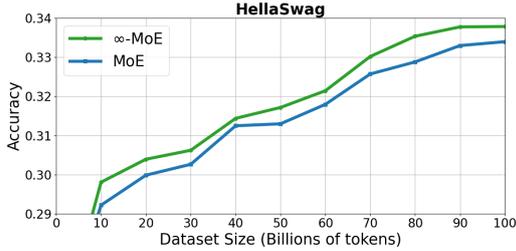


Figure 3: Accuracy on HellaSwag as a function of training data size (in billions of tokens).  $\infty$ -MoE is compared against a MoE baseline (GPT-2 Small backbone).

Table 2: Ablation on the dimension of the continuous index  $z$ . Lower is better.

$d_z$	PPL
32	3.244
64	3.220
128	<b>3.213</b>
256	3.219

#### 5.4 Routing Stability

We measured routing stability using the normalized routing entropy. For each mini-batch, we sampled  $z_t^{(k)} \sim p(z | x_t)$  ( $k = 1, \dots, K$ ) and obtained the active neuron set  $S(z_t^{(k)}) \subseteq \{1, \dots, d_{\text{ff}}\}$  from the top- $N\%$  mask. We defined the empirical selection distribution over FFN neurons as

$$q_j = \frac{\sum_{t,k} \mathbf{1}[j \in S(z_t^{(k)})]}{\sum_{t,k} |S(z_t^{(k)})|} \quad (j = 1, \dots, d_{\text{ff}}), \quad (8)$$

and computed

$$\tilde{H}(q) = \frac{-\sum_{j=1}^{d_{\text{ff}}} q_j \log q_j}{\log d_{\text{ff}}} \in [0, 1]. \quad (9)$$

Higher  $\tilde{H}(q)$  indicates more uniform utilization. As shown in Table 3,  $\tilde{H}(q)$  quickly stabilized at a high value during training.

Table 3: Normalized routing entropy  $\tilde{H}(q)$  during training.

Training Step	$\tilde{H}(q)$
5k	0.966
10k	0.902
15k	0.893
20k	0.893

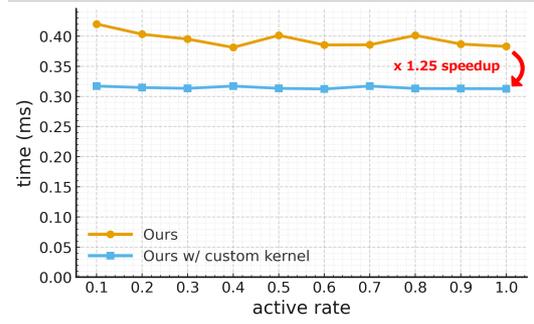


Figure 4: FFN forward latency vs. active rate. The custom kernel achieved 1.25-fold speedup.

#### 5.5 Latency Profiling: Custom Kernel

We measured end-to-end FFN forward latency (ms) while varying the active rate (percentage of FFN hidden units kept by the mask). All runs used the same batch size and sequence length on a single GPU after warmup. We compared a PyTorch reference (Ours) with a fused CUDA implementation that applied the mask and reduction in one pass (Ours w/ custom kernel). Results (Fig. 4) show consistently lower latency for the custom kernel and approximately 1.25-fold speedup; latency remained nearly flat across active rates for the custom kernel.

### 6 Conclusion

We present  $\infty$ -MoE, which extends MoE from a finite set of experts to a continuous (effectively infinite) expert space. It activates only a few sampled experts per token, preserving MoE-like efficiency while improving accuracy. On GPT-2 Small and Medium,  $\infty$ -MoE outperformed Switch and standard MoE, and tuning the number of samples  $K$  at inference provides a clear speed-accuracy tradeoff.

#### Limitations

Although  $\infty$ -MoE offers a promising framework for extending MoE models to an infinite expert space, several open challenges remain:

##### 1. Scaling Beyond GPT-2 Medium.

Although our experiments focused on GPT-2 Small and Medium, the behavior of  $\infty$ -MoE when scaling to larger models (e.g., GPT-3 and beyond) is not yet fully understood. In particular, it is unclear how performance and efficiency will change when:

- Increasing the total number of parameters while keeping the active (per-token)

parameter count fixed.

- Scaling both active and total parameters in tandem.

These scenarios raise questions about potential bottlenecks and tradeoffs in both training and inference at extreme scales.

## 2. Router Distributions.

Our current implementation employs a unimodal Gaussian router for simplicity. However, richer distributions, such as mixtures of Gaussians or nonparametric density estimators, could offer more expressive expert allocations, especially in high-dimensional expert spaces. Although this approach may improve coverage of diverse input patterns, designing efficient sampling and sparse-inference mechanisms becomes more complex, and variance reduction in training remains an open challenge.

## 3. Applicability to Other Domains.

Although our study highlights  $\infty$ -MoE’s utility in language modeling, it remains unclear how readily this framework generalizes to other domains such as vision (e.g., ViT) or multimodal vision-language models (VLMs). Practical concerns include adapting continuous expert indices to handle different input modalities, ensuring sparse and efficient routing for high-resolution data, and maintaining competitive accuracy in tasks beyond NLP.

4. **Training FLOPs and scalability.** From a compute perspective,  $\infty$ -MoE preserves the MoE property that per-token FFN FLOPs scale with the activated fraction of parameters. Theoretically, with  $K$  samples and an FFN active ratio  $r$  (top- $N\%$  masking), the dominant FFN compute is approximately  $\mathcal{O}(K \cdot r)$  of a dense FFN, and the routing overhead is small. However, lower FLOPs do not always translate to faster training in practice. Even with our fused CUDA kernel (Figure 4), end-to-end throughput can be slower than standard MoE systems because the mask depends on the input token. Since the set of active hidden units changes across tokens and samples, it is difficult to batch computation into fixed-shape GEMMs (general matrix multiplications, i.e., standard GPU matrix-multiply kernels) and to reuse a single highly optimized kernel. As

a result, naive implementations often rely on extra indexing and accumulation steps, which makes kernel fusion and vectorized execution harder and reduces practical hardware utilization.

In our experiments, we prioritized throughput over strictly minimizing FLOPs, and implemented the FFN computation using masked dense matrix multiplications that include the masking operation in the matmul pipeline. This design allows using fixed-shape GEMMs with shared weights across samples, enabling highly optimized kernels (e.g., Tensor Core) and improving wall-clock efficiency, at the cost of performing additional arithmetic compared to ideal sparse execution. Closing the remaining gap between theoretical compute savings and practical speed likely requires further system-level optimizations, as well as better compiler/runtime or hardware support for dynamic sparsity.

## References

Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Qin Cai, Vishrav Chaudhary, Dong Chen, Dongdong Chen, Weizhu Chen, Yen-Chun Chen, Yi-Ling Chen, Hao Cheng, Parul Chopra, Xiyang Dai, Matthew Dixon, Ronen Eldan, Victor Fragoso, Jianfeng Gao, Mei Gao, Min Gao, Amit Garg, Allie Del Giorno, Abhishek Goswami, Suriya Gunasekar, Emman Haider, Junheng Hao, Russell J. Hewett, Wenxiang Hu, Jamie Huynh, Dan Iter, Sam Ade Jacobs, Mojan Javaheripi, Xin Jin, Nikos Karampatziakis, Piero Kauffmann, Mahoud Khademi, Dongwoo Kim, Young Jin Kim, Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuanzhi Li, Yunsheng Li, Chen Liang, Lars Liden, Xihui Lin, Zeqi Lin, Ce Liu, Liyuan Liu, Mengchen Liu, Weishung Liu, Xiaodong Liu, Chong Luo, Piyush Madan, Ali Mahmoudzadeh, David Majercak, Matt Mazzola, Caio César Teodoro Mendes, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norick, Barun Patra, Daniel Perez-Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Liliang Ren, Gustavo de Rosa, Corby Rosset, Sambudha Roy, Olatunji Ruwase, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacrose, Shital Shah, Ning Shang, Hiteshi Sharma, Yelong Shen, Swadheen Shukla, Xia Song, Masahiro Tanaka, Andrea Tupini, Praneetha Vaddamanu, Chunyu Wang, Guanhua Wang, Lijuan Wang, Shuohang Wang, Xin Wang, Yu Wang, Rachel Ward, Wen Wen, Philipp Witte, Haiping Wu, Xiaoxia Wu, Michael Wyatt, Bin Xiao, Can Xu, Jiahang Xu,

- Weijian Xu, Jilong Xue, Sonali Yadav, Fan Yang, Jianwei Yang, Yifan Yang, Ziyi Yang, Donghan Yu, Lu Yuan, Chenruidong Zhang, Cyril Zhang, Jianwen Zhang, Li Lyna Zhang, Yi Zhang, Yue Zhang, Yunan Zhang, and Xiren Zhou. 2024. [Phi-3 technical report: A highly capable language model locally on your phone](#). *Preprint*, arXiv:2404.14219.
- Pratyay Banerjee, Kuntal Kumar Pal, Arindam Mitra, and Chitta Baral. 2019. [Careful selection of knowledge to solve open book question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6120–6129, Florence, Italy. Association for Computational Linguistics.
- Michael Boratko, Harshit Padigela, Divyendra Mikilineni, Pritish Yuvraj, Rajarshi Das, Andrew McCallum, Maria Chang, Achille Fokoue-Nkoutche, Pavan Kapanipathi, Nicholas Mattei, Ryan Musa, Kartik Talamadupula, and Michael Witbrock. 2018. [A systematic classification of knowledge, reasoning, and context within the ARC dataset](#). In *Proceedings of the Workshop on Machine Reading for Question Answering*, pages 60–70, Melbourne, Australia. Association for Computational Linguistics.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. [Evaluating large language models trained on code](#). *arXiv preprint arXiv:2107.03374*.
- Aidan Clark, Diego De Las Casas, Aurelia Guy, Arthur Mensch, Michela Paganini, Jordan Hoffmann, Bogdan Damoc, Blake Hechtman, Trevor Cai, Sebastian Borgeaud, George Bm Van Den Driessche, Eliza Rutherford, Tom Hennigan, Matthew J Johnson, Albin Cassirer, Chris Jones, Elena Buchatskaya, David Budden, Laurent Sifre, Simon Osindero, Oriol Vinyals, Marc’Aurelio Ranzato, Jack Rae, Erich Elsen, Koray Kavukcuoglu, and Karen Simonyan. 2022. [Unified scaling laws for routed language models](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 4057–4086. PMLR.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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. 2024. [Deepseekmoe: Towards ultimate expert specialization in mixture-of-experts language models](#). *Preprint*, arXiv:2401.06066.
- William Fedus, Barret Zoph, and Noam Shazeer. 2022. [Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity](#). *Journal of Machine Learning Research*, 23(120):1–39.
- Xu Owen He. 2024. [Mixture of a million experts](#).
- Robert Jacobs, Michael Jordan, Steven Nowlan, and Geoffrey Hinton. 1991. [Adaptive mixtures of local experts](#). *Neural Computation*, 3:79–87.
- 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, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th ophile Gervet, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2024. [Mixture of experts](#). *Preprint*, arXiv:2401.04088.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#). *CoRR*, abs/2001.08361.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. [RACE: Large-scale ReAding comprehension dataset from examinations](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark. Association for Computational Linguistics.
- Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2020. [Gshard: Scaling giant models with conditional computation and automatic sharding](#). *Preprint*, arXiv:2006.16668.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *arXiv preprint arXiv:2107.13586*.
- Jan Ludziejewski, Jakub Krajewski, Kamil Adamczewski, Maciej Pi ro, Micha  Krutul, Szymon Antoniak, Kamil Ciebiera, Krystian Kr l, Tomasz Odrzyg dz, Piotr Sankowski, Marek Cygan, and Sebastian Jaszczur. 2024. [Scaling laws for fine-grained mixture of experts](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 33270–33288. PMLR.
- Guilherme Penedo, Hynek Kydl ek, Loubna Ben al-lal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. 2024. [The fineweb datasets: Decanting the web for the finest text data at scale](#). In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). *OpenAI*. Accessed: 2024-11-15.

Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. [Winogrande: an adversarial winograd schema challenge at scale](#). *Commun. ACM*, 64(9):99–106.

Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. [Outrageously large neural networks: The sparsely-gated mixture-of-experts layer](#). *Preprint*, arXiv:1701.06538.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a machine really finish your sentence?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.

## A Hyperparameter

Details are provided in Table 4.

## B Total Computation for Experiments

We executed the experiments mainly by running the training for each model using eight nodes, each equipped with eight NVIDIA H200 (141GB) GPUs.

## C License

### C.1 Model

- GPT-2 small/medium: Modified MIT License

### C.2 Dataset

- FineWeb: Open Data Commons Attribution License (ODC-By) v1.0

## D Additional Experimental Results

### D.1 Scaling with dataset size

To complement subsection 5.2, Table 5 reports zero-shot accuracy on the full benchmark suite for GPT-2 Small-based models trained on 100B tokens. At this larger data scale,  $\infty$ -MoE achieves a higher average score (0.412 vs. 0.407) and improves 5 out of 7 tasks. These results indicate that the gains from continuous expert indexing persist beyond the 10B-token setting used in the main experiments.

### D.2 Comparison with More Experts

To test whether increasing the number of discrete experts is sufficient, we compare  $\infty$ -MoE with a 16-expert MoE baseline (top-2) trained on 10B tokens. We keep the active parameter budget the same by using top-2 activation in both models. As shown in Table 6,  $\infty$ -MoE achieves lower perplexity, suggesting better parameter efficiency than simply adding more discrete experts.

Table 6: Perplexity comparison with a 16-expert MoE baseline under a matched active-parameter budget (top-2, trained on 10B tokens). Lower is better.

Model	PPL
MoE (Top-2 over 16 experts)	3.294
Ours (Top-2 over infinite experts)	<b>3.238</b>

## E Use of Large Language Models

We employed a large language model (ChatGPT; “GPT-5 Thinking”) only for English-language polishing and light copy-editing. The model was not used to generate ideas and experimental design. All technical content and claims were authored and verified by the human authors. No non-public data, confidential information, or personally identifiable information were provided to the model.

Table 4: Model and training hyperparameters used in the experiments.

Parameter	GPT2-small	GPT2-medium
<b>Model Hyperparameters</b>		
Block size	1024	1024
Vocab size	50257	50257
Layers	12	24
Heads	12	16
Embedding dim	768	1024
Hidden dim	3072	4096
Gate dim(z dim)	256	256
<b>Training Hyperparameters</b>		
Total batch size	524288	
Gradient accumulation steps	1	
Optimizer	adamw	
Learning rate	0.0006	
Weight decay	0.1	
Warmup ratio	0.03	
Warmup iterations	700	
Data type	bfloat16	
ZeRO stage	1	

Table 5: Comparison of MoE and  $\infty$ -MoE (Ours) on the full benchmark suite. The models are GPT-2 Small based and trained on 100B tokens.

Model	BoolQ	HellaSwag	WinoGrande	ARC-e	ARC-c	OBQA	RACE-h	Avg
MoE	0.592	0.334	<b>0.511</b>	0.485	0.198	<b>0.190</b>	0.534	0.407
Ours ( $\infty$ -MoE)	<b>0.605</b>	<b>0.335</b>	0.510	<b>0.490</b>	<b>0.206</b>	0.170	<b>0.567</b>	<b>0.412</b>