

---

# Sparse Upcycling: Inference Inefficient Finetuning

---

**Sasha Doubov**

Databricks

sasha.doubov@databricks.com

**Nikhil Sardana**

Databricks

nikhil@databricks.com

**Vitaliy Chiley**

Databricks

vitaliy.chiley@databricks.com

## Abstract

Small, highly trained, open-source large language models are widely used due to their inference efficiency, but further improving their quality remains a challenge. Sparse upcycling is a promising approach that transforms a pretrained dense model into a Mixture-of-Experts (MoE) architecture, increasing the model’s parameter count and quality. In this work, we compare the effectiveness of sparse upcycling against continued pretraining (CPT) across different model sizes, compute budgets, and pretraining durations. Our experiments show that sparse upcycling can achieve better quality, with improvements of over 20% relative to CPT in certain scenarios. However, this comes with a significant inference cost, leading to 40% slowdowns in high-demand inference settings for larger models. Our findings highlight the trade-off between model quality and inference efficiency, offering insights for practitioners seeking to balance model quality and deployment constraints.

## 1 Introduction

Recent advancements have made several small, open-source large language models (LLMs) widely available to practitioners, such as Llama 3 8B [Dubey et al., 2024], Gemma 2B, and 9B models [Team et al., 2024]. These dense decoder-only models are designed for inference efficiency: smaller than their flagship counterparts (often 70B+ parameters), but trained on trillions of tokens for improved quality. Despite their high quality, there remains a persistent demand for further improvements, particularly in downstream task performance.

One common approach to improve model quality is continued pretraining (CPT), where the model is further trained on additional data, typically on a different dataset than the original pretraining data [Gupta et al., 2023]. While CPT can improve model quality, it is limited by the original model’s dense architecture and parameter count.

An alternative and increasingly popular approach is sparse upcycling [Komatsuzaki et al., 2023], which increases a model’s parameter count by converting a dense model into a Mixture-of-Experts (MoE) model. MoE architectures dynamically activate only a subset of weights (experts) during inference, allowing the model to expand its parameter count without proportionally scaling up inference costs [Shazeer et al., 2017]. Sparse upcycling a dense model into an MoE architecture has the potential to improve model quality, but it introduces a trade-off: sparse upcycled models are significantly larger than their dense counterparts, resulting in higher inference costs and limiting their utility for large-scale real-world deployments. Sparse upcycling thus conflicts with the growing trend to deploy smaller models specifically to decrease the cost of inference [Sardana et al., 2024].

In this work, we ask: **What is the trade-off between model quality and inference efficiency for sparse upcycling?** We compare sparse upcycling to dense CPT across varying model sizes, compute

budgets, and pretraining durations. By exploring these trade-offs, we provide insights for practitioners on how to balance performance gains with the practical costs of deploying LLMs.

## 2 Related Work

Mixture-of-Experts (MoE) [Shazeer et al., 2017, Fedus et al., 2022, Gale et al., 2022] is a sparse model architecture in which input tokens are dynamically routed to different weights (experts) in a model. The typical MoE language model architecture modifies the transformer block: the attention weights stay the same, but rather than having a single multi-layer perceptron (MLP) process tokens, a routing function is introduced which dynamically routes tokens to a subset of different MLPs. This form of dynamic sparsity leads to more compute-efficient training and better inference performance. Popular open-source MoEs have been released [Team, 2024a, Jiang et al., 2024], typically at larger model sizes than the smallest dense models available.

Sparse upcycling was introduced in Komatsuzaki et al. [2023], which experimented with upcycling for vision and language modeling tasks. The authors duplicated a transformer’s MLPs  $n$  times to create  $n$  MoE experts and replaced half of the MLP layers with MoE layers. They ablated many design choices, such as different routing mechanisms, adding Gaussian noise to the expert weights and pretraining duration. Unlike our setting, the authors studied encoder-decoder T5 models [Raffel et al., 2023], rather than the decoder-only auto-regressive models that are currently popular in the open-source community Komatsuzaki et al. [2023] also performed upcycling and CPT in a similar optimization setting as the pretraining phase: they use the *same* data as during pretraining, and the learning rate schedule is seamlessly continued due to the use of inverse-sqrt learning rate schedule. In practice, open-source models have been cosine annealed and the original data is not available [Dubey et al., 2024], making it difficult to apply existing results.

More recent works have more complex training recipes, by training  $n$  experts on  $n$  different datasets, later merging them into an MoE. Sukhbaatar et al. [2024] shows that doing data-specific dense finetunes, averaging the attention weights and creating different MLP experts is a more efficient training method than simply duplicating MLPs during upcycling. Zhang et al. [2024] replaces the averaging step for attention with parallel attention modules. However, Wei et al. [2024] has shown that the gap between data-specific upcycling and vanilla (duplicating MLPs) sparse upcycling diminishes over longer training durations. We chose to focus on the vanilla sparse upcycling recipe in this work with further discussion in Section 6. Past work has not benchmarked inference performance of upcycled models, a focus of our work.

## 3 Upcycling Quality Improvements

### 3.1 Experimental Setup

We have two distinct phases of training: dense pretraining and continued pretraining/upcycling. During the dense pretraining phase, we train a model on a generic common crawl, while CPT/upcycling is done with a higher quality domain data. Our dense pretrained models are fully cosine annealed, and we perform warmup during the CPT/upcycling training phase.

We experiment with two dense model sizes: a 436M and 1.4B model. Further experimental details are described in Appendix A.

Table 1: Pretraining models and training durations.

Model Size	Duration	Tokens
436M	Medium	43B
	Long	100B
	Extra Long	200B
1.4B	Medium	142B
	Long	354B

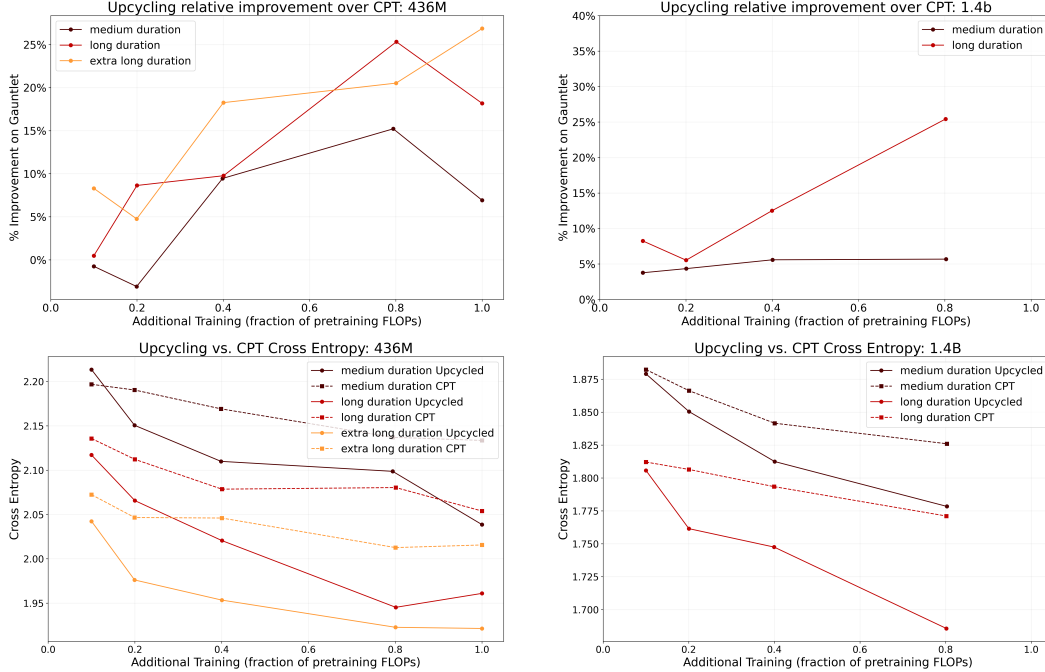


Figure 1: **Results from CPT vs. upcycling:** The top row shows the relative improvements on Gauntlet Core Average as a function of additional training. The bottom row compares the final cross entropy scores of the upcycled and CPT models.

### 3.1.1 Pretraining

We vary the amount of pretraining tokens for the different model sizes, as shown in Table 1. Note that we train our models for a large amount of pretraining tokens relative to the model size, with further discussion in Appendix A.3.

### 3.1.2 CPT/Upcycling

We create a Mixture-of-Experts from a dense checkpoint by duplicating the GLU layers within a block 8 times, with 8 experts and randomly initialize the router weights. We use Top- $K$  learned dropless routing following Gale et al. [2022] with  $K = 2$ . The resulting upcycled models contain 1.6B parameters (originally 436M) and 6.7B parameters (originally 1.4B). We perform different CPT/upcycling runs and match FLOP budgets for each run.

We report downstream results on the Eval Gauntlet v0.3 [MosaicML et al., 2023], a benchmark of 35 in-context learning tasks. We report the aggregate Gauntlet Core Average score, which is normalized and averaged across all of the tasks. More details are available in Appendix A.5.

## 3.2 Results

We show iso-FLOP comparisons of upcycling and CPT in Figure 1. In general, upcycling tends to achieve lower loss compared to the corresponding CPT run, showing that upcycling is indeed able to produce better quality models. For both CPT and upcycling, there are diminishing returns from extending training arbitrarily, although CPT saturates earlier than the upcycling runs.

Figure 1 also shows the relative improvements on Gauntlet Core Average across different models and durations. We find that a significant fraction of pretraining is needed to improve downstream performance — at 40% of the full training budget we see gains up to 20% with the 436M model and less than 15% with the 1.4B model.

Hence, upcycling is able to achieve lower loss and improve downstream quality, albeit it requires a significant portion of the original compute budget.

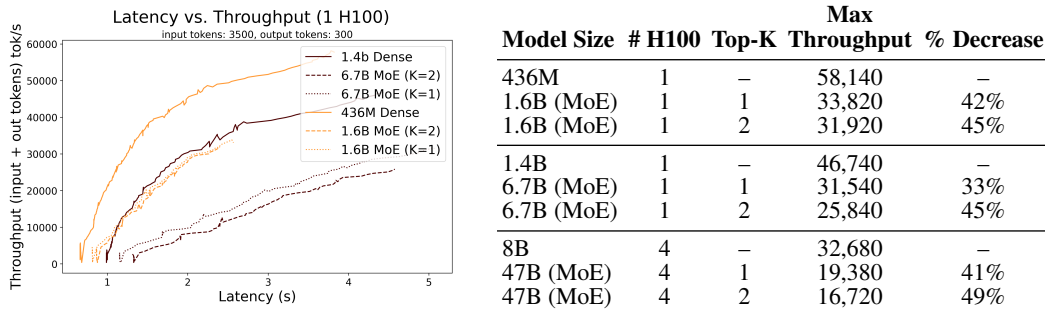


Figure 2: Inference Speed of CPT vs. Upcycled Models.

## 4 Inference Efficiency

### 4.1 Experimental Setup

To show the difference in inference efficiency between CPT vs. upcycled models, we benchmark latency and throughput of our 436M and 1.4B dense models along with their upcycled counterparts. In addition, we measure an 8B dense and corresponding 47B upcycled model to understand inference for upcycling a model at the popular 7–8B scale.

For the MoE models, we include benchmarking with top- $K=1$  and top- $K=2$ . This value sets the sparsity of the MoE model: a token is routed to  $K$  experts as part of the FFN computation. A higher top- $K$  value leads to more active parameters and correspondingly more FLOPs used during a model pass. When an upcycled model uses top- $K=1$ , it approximately matches the FLOPs of the original dense model, only with the relatively small overhead of router computation.

We use input: 3500 tokens, output: 300 tokens. We measure latency (the total time in seconds to process input + output tokens per request) and throughput (the amount of input + output tokens that the system can process per second). We use Nvidia’s H100 GPUs and the vLLM inference engine in our benchmarking setup. Further benchmarking details are available in Appendix B.

### 4.2 Results

Dense models outperform their upcycled counterparts significantly at inference time, as shown in Figure 2. In the high request regime, we compare the maximum throughput achieved by the dense models vs. upcycled models, and find significant decreases across the board. This is especially true for the larger 1.5B and 8B models, and even holds when top- $K=1$ .

In the high request regime, ie. when comparing max throughput, one would expect the performance for the dense model and the top- $K=1$  to be comparable. This is because the model is operating in the compute-bound regime, where the number of active parameters determines inference speed. However, we still see a large gap in maximum throughput between dense and MoE models. This may be due a few different factors. Dense models have less total parameters than their upcycled counterparts, which means that they use less GPU memory and can support higher batch sizes, allowing them to achieve higher throughput. As well, there is likely room for MoE-specific optimizations in vLLM that can further close the gap between dense and MoE model performance [Adnan Hoque, 2024].

As mentioned above, the increased parameter count in upcycled models leads them to using more GPU memory for storing the weights. The higher memory usage limits the hardware configurations suitable for running the upcycled models or necessitates model compression techniques.

Hence, our empirical results show a large gap in MoE and dense inference performance, even in regimes where it would be expected to have less of a performance gap.

## 5 Discussion

Larger upcycled models are empirically slower for inference, with maximum throughput reductions of 30-50%. Using  $K = 1$  leads to smaller throughput drops, although the gap remains significant.

We find that upcycling benefits from longer training durations, offering significant benefits over CPT after training for over 20-40% of the original pretraining budget. Hence, a significant amount of FLOPs must be sunk into upcycling to reap the benefits of the additional parameters. Upcycling is a good fit with a large FLOP budget and for applications where quality significantly outweighs inference performance.

Past work has not focused on the inference overhead of upcycling, emphasizing model quality. They also did not show real-world inference performance of upcycled models. However, given the growing trend of training smaller models for inference and inference-time compute [Snell et al., 2024], we believe that the inference cost of upcycling is an important consideration when trying to improve model quality.

## 6 Limitations

We note that our work has several limitations, in particular:

1. **Fixed MoE architecture:** We do not ablate the number of experts or top- $K$  value during training. We believe that 8 experts and  $K = 2$  is a reasonable architecture used in open-source models such as Mixtral [Jiang et al., 2024] and keeps the total parameters of the model relatively low. We benchmark top- $K=1$  to see how matching MoE and dense active parameters affects real-world inference performance.
2. **Inference Setup:** Our benchmarking setup uses vLLM as our inference engine. While this is a popular inference engine, we expect performance can differ significantly across different engines and that performance for MoEs can differ as well [Team, 2024b]. Still, we feel our work provides an idea of out-of-the-box performance of serving upcycled MoEs.
3. **Vanilla upcycling recipe:** We replicate the GLU layers rather than doing dense fine-tuning on different domains and then merging the models. While this more complex recipe may increase training efficiency gains, we choose to focus on the simpler setting and believe this is a promising technique for future work.
4. **Assumes pre-existing dense models:** Our work does not compare against training MoEs from scratch, a more FLOP efficient architecture in the pretraining setting [Fedus et al., 2022, Team, 2024a, Gale et al., 2022]. We instead focus solely on the post-training setting. This is due to the availability of high-quality open-source dense models that practitioners may wish to use to improve quality.

## 7 Conclusion

Our findings show that substantial additional training is required for an upcycled model to offset the extra inference overhead it introduces. While upcycling does improve model quality, it demands significant computational resources in terms of training FLOPs. Our empirical benchmarks reveal a considerable performance gap between upcycled models and dense models during inference.

Future research could explore alternative upcycling methods that address inference costs, as well as focus on further enhancing the inference performance of MoE models.

## References

- Antoni Virós Martin Chih-Chieh Yang Adnan Hoque, Less Wright. Accelerating moe model inference with locality-aware kernel design, 2024. URL <https://pytorch.org/blog/accelerating-moe-model/>.
- Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. Mathqa: Towards interpretable math word problem solving with operation-based formalisms. *CoRR*, abs/1905.13319, 2019. URL <http://arxiv.org/abs/1905.13319>.
- Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7432–7439, 2020.
- Cody Blakeney, Mansheej Paul, Brett W. Larsen, Sean Owen, and Jonathan Frankle. Does your data spark joy? performance gains from domain upsampling at the end of training, 2024. URL <https://arxiv.org/abs/2406.03476>.
- Christopher Clark, Kenton Lee, Tom Kwiatkowski Ming-Wei Chang, Michael Collins, and Kristina Toutanova. Boolq: Exploring the surprising difficulty of natural yes/no questions. In *NAACL*, 2019.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv:1803.05457v1*, 2018.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity, 2022. URL <https://arxiv.org/abs/2101.03961>.
- Trevor Gale, Deepak Narayanan, Cliff Young, and Matei Zaharia. Megablocks: Efficient sparse training with mixture-of-experts, 2022. URL <https://arxiv.org/abs/2211.15841>.
- Kshitij Gupta, Benjamin Thérien, Adam Ibrahim, Mats L. Richter, Quentin Anthony, Eugene Belilovsky, Irina Rish, and Timothée Lesort. Continual pre-training of large language models: How to (re)warm your model?, 2023. URL <https://arxiv.org/abs/2308.04014>.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. Training compute-optimal large language models, 2022. URL <https://arxiv.org/abs/2203.15556>.
- 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, Léo Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Théophile Gervet, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mixtral of experts, 2024. URL <https://arxiv.org/abs/2401.04088>.

- Aran Komatsuzaki, Joan Puigcerver, James Lee-Thorp, Carlos Riquelme Ruiz, Basil Mustafa, Joshua Ainslie, Yi Tay, Mostafa Dehghani, and Neil Houlsby. Sparse upcycling: Training mixture-of-experts from dense checkpoints, 2023. URL <https://arxiv.org/abs/2212.05055>.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.
- Hector Levesque, Ernest Davis, and Leora Morgenstern. The winograd schema challenge. In *Thirteenth International Conference on the Principles of Knowledge Representation and Reasoning*. Citeseer, 2012.
- Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. Logiqa: A challenge dataset for machine reading comprehension with logical reasoning. *CoRR*, abs/2007.08124, 2020. URL <https://arxiv.org/abs/2007.08124>.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *Conference on Empirical Methods in Natural Language Processing*, 2018. URL <https://api.semanticscholar.org/CorpusID:52183757>.
- MosaicML et al. Introducing mpt-7b: A new standard for open-source, commercially usable llms, 2023. URL [www.mosaicml.com/blog/mpt-7b](http://www.mosaicml.com/blog/mpt-7b). Accessed, pages 05–05, 2023.
- MosaicML NLP Team. Llm evaluation scores. <https://www.mosaicml.com/llm-evaluation>, 2023. Accessed: 2023-09-28.
- Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Quan Ngoc Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. The lambada dataset: Word prediction requiring a broad discourse context. *arXiv preprint arXiv:1606.06031*, 2016.
- Arkil Patel, Satwik Bhattamishra, and Navin Goyal. Are NLP models really able to solve simple math word problems? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2080–2094, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.168. URL <https://aclanthology.org/2021.naacl-main.168>.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer, 2023. URL <https://arxiv.org/abs/1910.10683>.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. SQuAD: 100,000+ Questions for Machine Comprehension of Text. *arXiv e-prints*, art. arXiv:1606.05250, 2016.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. CoQA: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266, 2019. doi: 10.1162/tacl\_a\_00266. URL <https://aclanthology.org/Q19-1016>.
- Melissa Roemmele, Cosmin Adrian Beja, and Andrew S Gordon. Choice of plausible alternatives: An evaluation of commonsense causal reasoning. *Papers from the 2011 AAAI Spring Symposium*, 2011.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. WINOGRANDE: an adversarial winograd schema challenge at scale. *CoRR*, abs/1907.10641, 2019. URL <http://arxiv.org/abs/1907.10641>.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan LeBras, and Yejin Choi. Socialiqa: Commonsense reasoning about social interactions. *arXiv preprint arXiv:1904.09728*, 2019.
- Nikhil Sardana, Jacob Portes, Sasha Doubov, and Jonathan Frankle. Beyond chinchilla-optimal: Accounting for inference in language model scaling laws, 2024. URL <https://arxiv.org/abs/2401.00448>.

- 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, 2017. URL <https://arxiv.org/abs/1701.06538>.
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters, 2024. URL <https://arxiv.org/abs/2408.03314>.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*, 2022.
- Sainbayar Sukhbaatar, Olga Golovneva, Vasu Sharma, Hu Xu, Xi Victoria Lin, Baptiste Rozière, Jacob Kahn, Daniel Li, Wen tau Yih, Jason Weston, and Xian Li. Branch-train-mix: Mixing expert llms into a mixture-of-experts llm, 2024. URL <https://arxiv.org/abs/2403.07816>.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. Commonsenseqa: A question answering challenge targeting commonsense knowledge. *arXiv preprint arXiv:1811.00937*, 2018.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024.
- Mosaic Research Team, Mar 2024a. URL <https://www.databricks.com/blog/introducing-dbrx-new-state-art-open-llm>.
- SGLang Team. Achieving faster open-source llama3 serving with sglang runtime (vs. tensorrt-llm, vllm), 2024b. URL <https://lmsys.org/blog/2024-07-25-sglang-llama3/>.
- Tianwen Wei, Bo Zhu, Liang Zhao, Cheng Cheng, Biye Li, Weiwei Lü, Peng Cheng, Jianhao Zhang, Xiaoyu Zhang, Liang Zeng, Xiaokun Wang, Yutuan Ma, Rui Hu, Shuicheng Yan, Han Fang, and Yahui Zhou. Skywork-moe: A deep dive into training techniques for mixture-of-experts language models, 2024. URL <https://arxiv.org/abs/2406.06563>.
- Thom Wolfe, Lewis Tunstall, and Patrick von Platen. Jeopardy dataset on hugging face hub. <https://huggingface.co/datasets/jeopardy>, 2022.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.
- Qizhen Zhang, Nikolas Gritsch, Dwaraknath Gnaneshwar, Simon Guo, David Cairuz, Bharat Venkitesh, Jakob Foerster, Phil Blunsom, Sebastian Ruder, Ahmet Ustun, and Acyr Locatelli. Bam! just like that: Simple and efficient parameter upcycling for mixture of experts, 2024. URL <https://arxiv.org/abs/2408.08274>.
- Wanjuan Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. Agieval: A human-centric benchmark for evaluating foundation models, 2023.

## A Experimental Details

### A.1 Data

For the dense pretraining phase, we use a generic common crawl data mix. For the continued pretraining/upcycling phase, we follow Blakeney et al. [2024] and use a higher quality mix of 4 broad categories with different proportions: Large-Scale Common Crawl (15%), code (35%), Small-Scale Common Crawl (15%), Domain Specific data (35%). The choice to have a different source of data between pretraining and CPT/upcycling is meant to emulate the setting of upcycling an open-source model, where the original dataset is rarely accessible .



## A.2 Models

We use the following hyperparameters described in Table 3. We performed a learning rate sweep at the shortest duration for CPT and upcycling with the *medium* duration dense runs, and use that learning rate for the rest of the experiments.

## A.3 Pretraining Duration

To select our pretraining duration, we chose to focus on the “over-trained” setting. The seminal “Chinchilla” work [Hoffmann et al., 2022] established scaling laws for optimally pretraining LLMs for a given training FLOP budget – finding an approximately constant optimal ratio of 20 training tokens per model parameter. In practice, models have been trained well-past this ratio of tokens/parameters, as smaller performant models are better for inference [Sardana et al., 2024]. For example, the Llama 3 8B and 70B models were pretrained for 15 trillion tokens, resulting in a token/parameter ratios of 1875 and 214 respectively [Dubey et al., 2024]. We include the token/parameters ratios for our model configurations in Table 2.

Table 2: Pretraining models and training durations with Token/Parameter Ratio.

Model Size	Duration	Tokens	Tokens/Parameter
436M	Medium	43.6B	100
	Long	100B	230
	Extra Long	200B	460
1.4B	Medium	142B	100
	Long	354B	250

## A.4 CPT/Upcycling Training Duration

We describe how much additional training was done for the corresponding models in Table 4.

## A.5 Evaluation Details

We report both the smoothed final cross entropy loss and downstream Evaluation Gauntlet scores for each training run. Our reported cross entropy loss is averaged over the final 50 steps of training to account for minor fluctuations.

We evaluate our model against version 3 of the open source Evaluation Gauntlet [MosaicML NLP Team, 2023], with tasks in five categories:

- **Commonsense Reasoning:** BIG-bench Strategy QA, BIG-bench Strange Stories [Srivastava et al., 2022], Common Sense QA [Talmor et al., 2018], COPA [Roemmele et al., 2011], OpenBook QA [Mihaylov et al., 2018], PIQA [Bisk et al., 2020], and SIQA [Sap et al., 2019].
- **Language Understanding:** LAMBADA [Paperno et al., 2016], HellaSwag [Zellers et al., 2019], Winograd Schema Challenge [Levesque et al., 2012], and Winogrande [Sakaguchi et al., 2019].
- **Reading Comprehension:** AGI Eval [Zhong et al., 2023], BoolQ [Clark et al., 2019], CoQA [Reddy et al., 2019], and SQuAD [Rajpurkar et al., 2016].
- **Symbolic Problem Solving:** AGI Eval SAT Math, AGI Eval LSAT [Zhong et al., 2023], BIG-bench Dyck Languages [Srivastava et al., 2022], BIG-bench Elementary Math QA [Srivastava et al., 2022], BIG-bench Operators [Srivastava et al., 2022], GSM8k [Cobbe et al., 2021], LogiQA [Liu et al., 2020], Math QA [Amini et al., 2019], and SVAMP [Patel et al., 2021].
- **World Knowledge:** ARC Easy, ARC Challenge [Clark et al., 2018], BIG-bench WikiData [Srivastava et al., 2022], Jeopardy [Wolfe et al., 2022], and MMLU [Hendrycks et al., 2020].

The Gauntlet Core Average is a simple average accuracy over all tasks. Each task is equally weighted after subtracting out the task’s baseline random accuracy and normalizing.

## B Inference Details

We use the vLLM v0.6.2 Kwon et al. [2023] inference engine for all inference benchmarks. For fair comparisons across MoE and dense vLLM implementations, we compare the performance of the Mistral and Mixtral

Table 3: 430M settings

<b>Model</b>	430M	1.4B
d_model	1024	2048
ffn hidden size	2560	5120
layers	22	24
heads	16	–
kv heads (GQA)	4	–
norm type	rmsnorm	–
tokenizer	gpt-4	–
total parameters	436M	1.42 M
<b>Pretraining Settings</b>		
optimizer	Lion	–
lr	8e-4	4e-4
weight decay	0.5	–
batch size	1024	–
seq len	4096	–
lr scheduler	cosine	–
warmup	1000500000 tokens	–
<b>CPT Settings</b>		
lr	5e-5	–
weight decay	0.05	–
batch size	1024	–
seq len	4096	–
lr scheduler	cosine	–
warmup	5e8 tokens	–
<b>Upcycling Settings</b>		
lr	1e-4	5e-5
weight decay	0.05	–
batch size	1024	–
seq len	4096	–
lr scheduler	cosine	–
warmup	5e8 tokens	–
load balancing loss	0.01	–
total parameters	1.6B	6.7B
active parameters	500M	1.9B

architectures [Jiang et al., 2024]. We also switch to the Llama 3 tokenizer for ease of use with vLLM and for consistency with the Llama3 8B model. We vary the Request Per Seconds at 0.1 increments in order to obtain the curves for Latency vs Throughput.

For the 430M and 1.4B models and their upcycled counterparts, we run inference on a single H100 GPU. For the 8B model and its upcycled counterpart we use 4 H100 GPUs and use tensor parallelism (TP). Note that we could have fit the 8B model *without* TP on a single H100, but chose to use the same number of gpus for both the dense and upcycled versions.

We average our results over 3 runs.

## C Additional Results

We show plots with absolute Gauntlet scores in Figure 4. We find that CPT models seem to saturate dramatically as compared to the CE loss shown in Figure 1, which is an interesting disconnect that requires more investigation.




Figure 3: Latency vs. Throughput - 8B vs 47B MoE

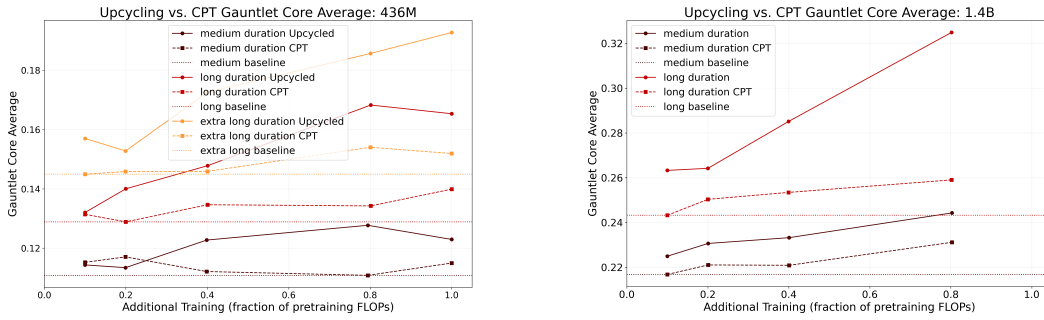


Figure 4: Absolute Gauntlet Scores for Dense vs. Upcycled Models.

Table 4: CPT Model sizes and additional training amounts

<b>Model Size</b>	<b>Duration</b>	<b>Additional Tokens</b>
436M	Medium	4.3B
		8.7B
		17.5B
		34.9B
		43.6B
436M	Long	10B
		20B
		40B
		80B
		100B
436M	Extra Long	20B
		40B
		80B
		100B
		200B
1.6B (Upycled)	Medium	3.2B
		6.5B
		13B
		26B
		32B
1.6B (Upycled)	Long	7.5B
		15B
		30B
		60B
		75B
1.6B (Upycled)	Extra Long	15B
		30B
		60B
		120B
		150B
1.4B	Medium	14B
		28B
		56B
		113B
1.4B	Long	35B
		70B
		142B
		284B
6.7B (Upycled)	Medium	9.6B
		19.3B
		38.6B
		77.2B
6.7B (Upycled)	Long	24B
		48B
		96B
		193B