


OLMoE: Open Mixture-of-Experts Language Models

Niklas Muennighoff^{ca} Luca Soldaini^a Dirk Groeneveld^a Kyle Lo^a Jacob Morrison^a
 Sewon Min^a Weijia Shi^w Pete Walsh^a Oyvind Tafjord^a Nathan Lambert^a
 Yuling Gu^a Shane Arora^a Akshita Bhagia^a Dustin Schwenk^a David Wadden^a
 Alexander Wettig^{ap} Binyuan Hui Tim Dettmers^a Douwe Kiela^c Ali Farhadi^{aw}
 Noah A. Smith^{aw} Pang Wei Koh^{aw} Amanpreet Singh^c Hannaneh Hajishirzi^{aw}
^aAllen Institute for AI ^cContextual AI ^wUniversity of Washington ^pPrinceton University
n.muennighoff@gmail.com hannah@allenai.org

Abstract

We introduce **OLMoE**, a fully open, state-of-the-art language model leveraging sparse Mixture-of-Experts (MoE). **OLMoE-1B-7B** has 7 billion (B) parameters but uses only 1B per input token. We pretrain it on 5 trillion tokens and further adapt it to create **OLMoE-1B-7B-INSTRUCT**. Our models outperform all available models with similar active parameters, even surpassing larger ones like Llama2-13B-Chat and DeepSeekMoE-16B. We present novel findings on MoE training, define and analyze new routing properties showing high specialization in our model, and open-source all our work: model weights, training data, code, and logs.

 **Weights** <https://hf.co/allenai/OLMoE-1B-7B-0924>
 **Data** <https://hf.co/datasets/allenai/OLMoE-mix-0924>
 **Code** <https://github.com/allenai/OLMoE>
 **Logs** <https://wandb.ai/ai2-llm/olmoe/reports/OLMoE-1B-7B-0924--Vmlldzo40TcyMjU3>




Figure 1: **Performance, cost, and degree of openness of open MoE and dense LMs.** Model names contain rounded parameter counts: model-active-total for MoEs and model-total for dense LMs. #ckpts is the number of intermediate checkpoints available. We highlight MMLU as a summary of overall performance; see §3 for more results. **OLMoE-1B-7B** performs best among models with similar active parameter counts and is the most open MoE.

1 Introduction

Despite significant advances in Large Language Models (LMs) on various tasks, there remains a clear trade-off between performance and cost in both training and inference. High-performing LMs are inaccessible for many academics and open-source developers as they are prohibitively expensive to build and deploy.¹ One approach to improve the cost-performance trade-off lies in using sparsely-activated Mixture-of-Experts (MoEs) [152]. MoEs have several experts in each layer, only a subset of which is activated at a time (see Figure C1). This makes MoEs significantly more efficient than dense models with a similar number of total parameters, which activate all parameters for every input [203]. For this reason, industry frontier models use MoEs including Gemini-1.5 [173] and reportedly GPT-4 [29].

Most MoE models, however, are closed-source: while some have publicly released model weights [43, 79, 156, 176, 178], they offer limited to no information about their training data, code, or recipes (see Figure 1). While there have been prior efforts to make language modeling research fully accessible [18, 65, 89, 102, 191, 207], they have been largely limited to dense LMs. This comes despite MoEs requiring *more* openness as they add complex new design questions to LMs, such as how many total versus active parameters to use, whether to use many small or few large experts, if experts should be shared, and what routing algorithm to use. The lack of open resources and findings about these details prevents the field from building cost-efficient open MoEs that approach the capabilities of closed-source frontier models.

To address these issues, we introduce **OLMOE**, a fully open Mixture-of-Experts language model with state-of-the-art performance among similarly-sized models. In particular, we pretrain **OLMOE-1B-7B** for 5.1 trillion tokens with 6.9B total parameters, of which only 1.3B are activated for each input token. This leads to a similar inference cost as using dense models with around 1B parameters, such as OLMo 1B [65] or TinyLlama 1B [208], but requires more GPU memory to store its 7B total parameters. Our experiments show that MoEs train $\sim 2\times$ faster than dense LMs with equivalent active parameters (Figure B1). In Figure 1, we show that **OLMOE-1B-7B** significantly outperforms all open 1B models and displays competitive performance to dense models with significantly higher inference costs and memory storage (e.g., similar MMLU scores to Llama2-13B, which is $\sim 10\times$ more costly). Via instruction- and preference tuning, we create **OLMOE-1B-7B-INSTRUCT**, which we find exceeds various larger instruct models including Llama2-13B-Chat [181], OLMo-7B-Instruct (0724), and DeepSeekMoE-16B [42] on common benchmarks (MMLU, GSM8k, HumanEval, etc.).

Our comprehensive set of controlled experiments highlights key design choices for MoEs (see Table 1) and LMs in general. One critical design decision for making MoEs performant is using fine-grained routing with granular experts [42]: we employ 64 small experts in each layer with 8 being activated. The choice of routing algorithm is also important: we find dropless [58] token-based routing [152] outperforms expert-based routing [217]. Our findings also include those that challenge prior work, such as the ineffectiveness of shared experts [42] and the limited benefits of sparsely upcycling a pretrained dense LM into an MoE [85] unless under small compute budgets. Finally, we present novel ways to analyze routing behavior in Mixture-of-Experts finding that for **OLMOE-1B-7B** routing saturates early in pretraining, experts are rarely co-activated, and experts exhibit domain and vocabulary specialization. We intend our fully open MoE to facilitate more research and analysis to improve our understanding of these models. We release training code, intermediate checkpoints (every 5000 steps), training logs, and training data under open-source licenses (Apache 2.0 <http://www.apache.org/licenses/LICENSE-2.0> or ODC-By 1.0 <https://opendatacommons.org/licenses/by/1-0/>).

2 Pretraining and adaptation

Pretraining **OLMOE** is a decoder-only LM consisting of N_L transformer [183] layers. The feedforward network (FFN) in dense models like OLMo [65], is replaced with an MoE module consisting of N_E smaller FFN modules called experts, of which a subset of k experts is activated for

¹For example, even with 16 H100 GPUs and several optimizations, Llama 3 405B only achieves a decoding throughput of around 100 tokens per second [50].

Table 1: **Key MoE design choices and our setup for OLMoE-1B-7B based on our experiments.** Full configuration for **OLMoE-1B-7B** is in §C.

Design choice	Description	Experiment	OLMoE-1B-7B
Active params	# active parameters per input token	§B.1.1	1.3B active
Total params	Total # of parameters in the model	§B.1.1	6.9B total
Expert granularity	Using fine-grained small experts vs. a few large experts [39]	§B.1.2	64 small experts with 8 activated
Expert sharing	Whether or not to include a shared expert [39]	§B.1.3	No shared expert
Routing algorithm	How inputs are assigned to experts, e.g., assignment on a per token basis (e.g., 2 experts per token) or per expert basis (e.g., 2 tokens per expert), and whether or not all tokens get assigned or some get dropped [58, 217]	§B.1.4	Dropless [58] MoE with token choice
Sparse upcycling	Whether to start from a dense model [85, 209]	§B.1.5	Not used
Load balancing loss	Auxiliary loss to penalize unequal assignment to experts that may harm performance [152]	§B.1.6	Used with weight 0.01
Router z-loss	Auxiliary loss to penalize large logits in the router that may cause instabilities [219]	§B.1.7	Used with weight 0.001

each processed input token x (also see Figure C1):

$$\text{MoE module}(x) = \sum_{i \in \text{Top-}k(r(x))} \text{softmax}(r(x))_i E_i(x) \quad (1)$$

where r , called the router, is a learned linear layer mapping from the input logits to the chosen k experts. A softmax is applied to the router outputs to compute routing probabilities for all N_E experts. Each selected expert E_i processes the input x , the output of which is then multiplied with its respective routing probability. The results are then summed across all chosen Top- k experts to constitute the output of the MoE module for a single layer of the model out of its N_L total layers. Key decisions in designing an MoE model include determining the number of activated and total parameters, the design of the experts (e.g., granularity, whether or not to include shared experts), and the choice of the routing algorithm. Moreover, training an MoE model can involve initializing from a dense model (sparse upcycling) and changing the training objective, such as including auxiliary load balancing and router z-losses. We run experiments to investigate each of these design choices in isolation in §B.1 and summarize our final decisions in Table 1. In summary, we use 1.3B active parameters out of a total of 6.9B, with 8 activated experts out of 64 per layer. We use dropless token choice routing [58]: For each input token, the learned router network determines 8 experts to process it. We train **OLMoE-1B-7B** from scratch with two auxiliary losses: load balancing loss (\mathcal{L}_{LB}) [152] and router z-loss (\mathcal{L}_{RZ}) [219], which we define and experiment with in §B.1.6 and §B.1.7, respectively. We multiply them with respective loss weights, α and β , and sum them linearly with the cross entropy loss (\mathcal{L}_{CE}) to arrive at our final training loss:

$$\mathcal{L} = \mathcal{L}_{CE} + \alpha \mathcal{L}_{LB} + \beta \mathcal{L}_{RZ} \quad (2)$$

For our pretraining data, we mix data from DCLM [89] and Dolma 1.7 [161], which includes: (1) a quality-filtered subset of Common Crawl, referred to as DCLM-Baseline, (2) StarCoder, Algebraic Stack and arXiv, used in both DCLM and Dolma 1.7, and (3) peS2o and Wikipedia from Dolma 1.7. We refer to our pretraining dataset as **OLMoE-MIX**. We train for a total of 5.133T tokens (1.3 epochs following Muennighoff et al. [120]) and provide data statistics in Table C1. Our full pretraining configuration for **OLMoE-1B-7B** is in §C.

Adaptation We create **OLMoE-1B-7B-INSTRUCT** by following a standard adaptation recipe split into **instruction tuning** [117, 188, 147, 154, 204] followed by **preference tuning** [31, 15, 136, 54] building on prior open models [182, 76, 186]. In our instruction tuning dataset, we add more code and math data to boost performance on downstream coding and math applications. Other models,

such as GPT-4 [126] and Llama 3 [50] similarly include samples from math datasets like GSM8k [35] or MATH [71] during pretraining. We also include No Robots and a subset of Daring Anteater as they are of high quality and add diversity, two key factors for successful adaptation [186, 214, 103, 119]. We describe our adaptation datasets in Table C2 and hyperparameters in §C.

3 Results

Our evaluation procedure consists of three parts: **During pretraining** (§F), **After pretraining**, and **After adaptation**. We detail the setup for each in §D.

Table 2: **OLMoE-1B-7B after pretraining versus larger MoEs and dense LMs.** We compare with dense LMs close to **OLMoE-1B-7B** either in active parameters (1B, approximates speed and cost) or total parameters (7B, approximates memory requirements). Model names contain rounded parameter counts: `model-active-total` for MoEs and `model-total` for dense LMs (this leads to some differences to official names, e.g., while called “Gemma2-2B” it actually has 2.6B active and total parameters [175]). Chall. = Challenge. We run all evaluations ourselves with 5 few-shots, see §D for details.

	Active params	Open Data	MMLU	Hella- Swag	ARC- Chall.	ARC- Easy	PIQA	Wino- Grande
LMs with ~7-9B active parameters								
Llama2-7B [181]	6.7B	✗	46.2	78.9	54.2	84.0	77.5	71.7
OLMo-7B (0724) [65]	6.9B	✓	54.9	80.5	68.0	85.7	79.3	73.2
Mistral-7B [78]	7.3B	✗	64.0	83.0	78.6	90.8	82.8	77.9
DCLM-7B [89]	6.9B	✓	64.4	82.3	79.8	92.3	80.1	77.3
Llama3.1-8B [50]	8.0B	✗	66.9	81.6	79.5	91.7	81.1	76.6
Gemma2-9B [175]	9.2B	✗	70.6	87.3	89.5	95.5	86.1	78.8
LMs with ~2-3B active parameters								
OpenMoE-3B-9B [197]	2.9B	✓	27.4	44.4	29.3	50.6	63.3	51.9
StableLM-2B [16]	1.6B	✗	40.4	70.3	50.6	75.3	75.6	65.8
DeepSeek-3B-16B [39]	2.9B	✗	45.5	80.4	53.4	82.7	80.1	73.2
JetMoE-2B-9B [156]	2.2B	✗	49.1	81.7	61.4	81.9	80.3	70.7
Gemma2-3B [175]	2.6B	✗	53.3	74.6	67.5	84.3	78.5	71.8
Qwen1.5-3B-14B [178]	2.7B	✗	62.4	80.0	77.4	91.6	81.0	72.3
LMs with ~1B active parameters								
Pythia-1B [18]	1.1B	✓	31.1	48.0	31.4	63.4	68.9	52.7
OLMo-1B (0724) [65]	1.3B	✓	32.1	67.5	36.4	53.5	74.0	62.9
TinyLlama-1B [208]	1.1B	✓	33.6	60.8	38.1	69.5	71.7	60.1
Llama3.2-1B [50]	1.2B	✗	38.2	67.3	43.5	71.6	73.7	62.5
DCLM-1B [89]	1.4B	✓	48.5	75.1	57.6	79.5	76.6	68.1
OLMoE-1B-7B	1.3B	✓	54.1	80.0	62.1	84.2	79.8	70.2

After pretraining In Table 2 we find that **OLMoE-1B-7B** performs best among models that use less than 2B active parameters, making it the cheapest option for many use cases of LMs. For larger budgets, Qwen1.5-3B-14B has stronger performance but more than double the active and total parameters of **OLMoE-1B-7B**. Despite using ~6–7× less compute per forward pass, **OLMoE-1B-7B** exceeds some dense LMs with 7B parameters like Llama2-7B [181], but not others like Llama3.1-8B [50]. Figure 1 compares MMLU with active parameters, a proxy for the value of a model given its cost. **OLMoE-1B-7B** is the state of the art in its cost regime.

After adaptation In Table 3, we benchmark our instruction (SFT) and preference (DPO) tuning of **OLMoE-1B-7B**. SFT improves our model on all tasks measured. We observe a >10× gain on GSM8k, likely due to our inclusion of additional math data to account for the relatively small

Table 3: **OLMoE-1B-7B after adaptation vs. other models.** The JetMoE chat model (<https://hf.co/jetmoe/jetmoe-8b-chat>) is broken leading to random scores thus we exclude it. Model names contain rounded parameter counts: model-active-total for MoEs and model-total for dense LMs. We run all evaluations ourselves (§D). Models use different mixes for adaptation, e.g., **OLMoE** is trained on an improved version of the pipeline used for OLMo models.

Task (→) Setup (→) Metric (→)	MMLU	GSM8k	BBH	Human-Eval	Alpaca-Eval 1.0	XSTest	IFEval	Avg
	0-shot EM	8-shot CoT EM	3-shot EM	0-shot Pass@10	0-shot %win	0-shot F1	0-shot Loose Acc	
OLMo-1B (0724)	25.0	7.0	22.5	16.0	-	67.6	20.5	-
+SFT	36.0	12.5	27.2	21.2	41.5	81.9	26.1	35.9
+DPO	36.7	12.5	30.6	22.0	50.9	79.8	24.2	37.4
OLMo-7B (0724)	50.8	32.5	36.9	32.3	-	80.8	19.6	-
+SFT	54.2	25.0	35.7	38.5	70.9	86.1	39.7	49.3
+DPO	52.8	9.0	16.6	35.0	83.5	87.5	37.9	49.1
JetMoE-2B-9B	45.6	43.0	37.2	54.6	-	68.2	20.0	-
+SFT	46.1	53.5	35.6	64.8	69.3	55.6	30.5	50.4
DeepSeek-3B-16B	37.7	18.5	39.4	48.3	-	65.9	13.5	-
+Chat	48.5	46.5	40.8	70.1	74.8	85.6	32.3	57.0
Qwen1.5-3B-14B	60.4	13.5	27.2	60.2	-	73.4	20.9	-
+Chat	58.9	55.5	21.3	59.7	83.9	85.6	36.2	57.3
OLMoE-1B-7B	49.8	3.0	33.6	22.4	-	59.7	16.6	-
+SFT	51.4	40.5	38.0	51.6	69.2	84.1	43.3	54.0
+DPO	51.9	45.5	37.0	54.8	84.0	82.6	48.1	57.7

amounts of math data during pretraining (§2). Our DPO model, which we refer to as **OLMoE-1B-7B-INSTRUCT**, has the highest average among all models benchmarked. It outperforms the chat version of Qwen1.5-3B-14B despite Qwen having $>2\times$ more parameters and its pretrained model outperforming **OLMoE-1B-7B** in Table 2. The 84% score on AlpacaEval exceeds much larger dense models on the leaderboard,² like Llama2-13B-Chat [181].

4 MoE analysis

By advancing open and cost-efficient models (§1), **OLMoE-1B-7B** enables new research into LMs and MoEs. Making use of our released intermediate checkpoints, data, and code, we define and analyze four properties specific to MoEs: Router saturation (§4.1), Expert co-activation (§4.2), Domain specialization (§4.3), and Vocabulary specialization (§4.4).

4.1 Router saturation

Router saturation, as a function of time t , represents the proportion of overlapping activated experts between the final checkpoint and an intermediary checkpoint at time t . Router saturation thus corresponds to whether the router weights are still learning which expert will process certain data. A value of 100% indicates that the router at the intermediate checkpoint will route to the same experts as the final checkpoint router. See §H.1 for the detailed formula used to calculate the value.

Figure 2 shows that, after 1% of pretraining, up to $\sim 60\%$ of routing to the top-8 activated experts has already saturated (right). Thus the model already uses the same 8 experts for given input data as it will at the end of pretraining. This early saturation aligns with prior work [197]. At 40% of pretraining, saturation reaches up to $\sim 80\%$. However, which top-1 expert has the highest routing probability saturates slower (left). We find that routing in later layers saturates earlier during pretraining. Layer 0 is an outlier saturating significantly more slowly than other layers. Dai et al. [39] do not use an MoE in the first layer as they find that load balancing converges more slowly for the first layer. This is likely linked to our findings on saturation. Because routing in the first layer saturates slower, the

²https://tatsu-lab.github.io/alpaca_eval/




Figure 2: **Router saturation during pretraining measured on a random 0.5% of the C4 validation data.** We compute saturation by comparing the routing to the top- k experts at four intermediate checkpoints (1, 10, 20, and 40% of pretraining) to the final pretraining checkpoint (Equation 5).

experts that certain input data get routed to frequently change. These changes may lead to one expert suddenly getting significantly more data than others thereby impairing load balancing. We are excited about future work further investigating what happens in the first layer by building on our open release.

4.2 Expert co-activation

We define expert co-activation as the proportion of times two specific experts, E_i and E_j , are simultaneously activated out of the total number of activations of one of them. A co-activation of 100% indicates that if E_i is activated, E_j is also always activated. A value of 0% indicates that the experts never co-occur. See §H.1 for the formula used to calculate the value.

Figure 3 shows that there is no strong co-activation among experts in layer 7, with only few exceptions. This may indicate that there is little redundancy across different experts. Layers 7 and 15 (Figure H1) show similar co-activation patterns with several groups of 3 or 2 experts that tend to get activated together. We investigate tokens that activate these experts in §4.4. Further, in §H.2 (Figure H7), we investigate whether experts across layers, rather than within one layer, tend to process tokens together.

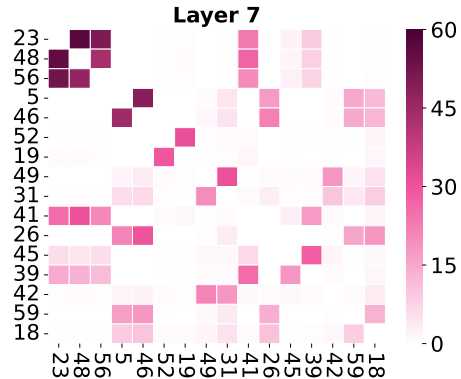


Figure 3: **Co-activation among experts of OLMoE-1B-7B on a random 0.5% of the C4 validation data.** We display the 32 experts with the highest maximum co-activation score via their expert IDs on the x- and y-axis. See Figure H1 for Layer 0 and 15.

4.3 Domain specialization

Domain specialization refers to the specialization of expert E_i to domain D , defined as the proportion of tokens from a particular domain D that get routed to a particular expert E_i (see §H.1 for the formula). A value of 100% indicates that all data from that domain is routed to E_i , whereas 0% indicates the expert is never used for that domain and can be removed from the model without affecting performance in that domain.

Figure 4 (left) shows many examples of experts that are activated significantly above or below random chance for *specific domains*. E.g., for arXiv, which has a very specific distribution with lots of scientific text, the first expert in layer 0 is nearly 100% specialized. This suggests that there is little redundancy in the knowledge of the experts in OLMoE-1B-7B, as they specialize in different kinds of data. GitHub and arXiv are often activated together in layer 7, which we explore further in §4.4. For *generic domains*, such as C4 [137], which is a web crawl containing various kinds of data, expert activations in OLMoE-1B-7B are much more balanced. This highlights that the load balancing (§B.1.6) works as intended and the model makes proper use of all experts for generic data. Mixtral-8x7B [79] in Figure 4 (right), however, exhibits little domain specialization across both

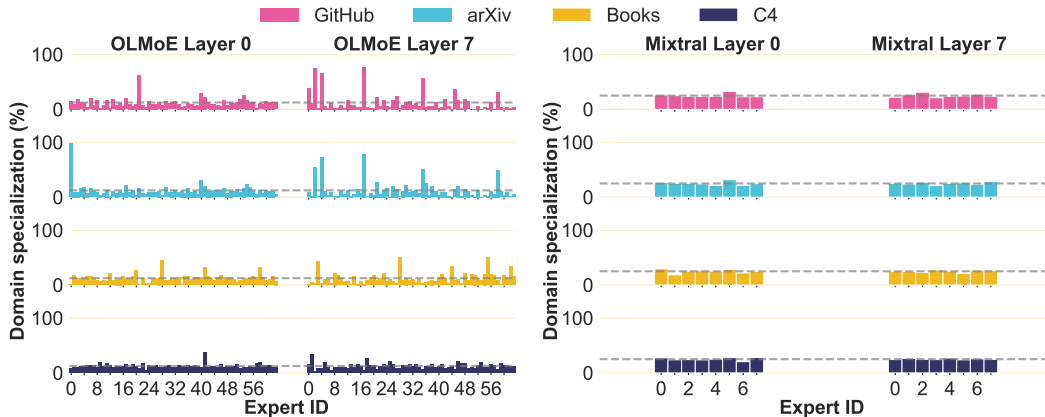


Figure 4: **Domain specialization of OLMoE-1B-7B (left) vs. Mixtral-8x7B (right)**. We visualize how often tokens from different domains get routed to the 64 (OLMoE) or 8 (Mixtral) experts after pretraining. We consider tokens routed to any of the $k = 8$ (OLMoE) or $k = 2$ (Mixtral) active experts (Equation 7). Horizontal gray lines are random chance or uniform routing ($8/64=12.5\%$ per expert for OLMoE-1B-7B with 8 active out of 64 total experts per layer and $2/8=25\%$ for Mixtral with 2 active out of 8 total experts per layer). See Figure H6 for $k = 1$ results.

unique and *generic* domains. Experts are activated close to the uniform routing baseline for all layers and domains. Thus, there may be more redundancy across experts in Mixtral, as they likely contain similar knowledge. We hypothesize that this is due to Mixtral being upcycled from Mistral [25]. The initialization from a dense model may limit the amount of possible specialization in the experts as they all start from the same local optimum. This is likely why training from scratch eventually outperforms upcycling in our pretraining experiments (§B.1.5).

4.4 Vocabulary specialization




Figure 5: **Vocabulary specialization of OLMoE-1B-7B across layers and experts**. To compute vocabulary specialization per layer (left) we average the specialization of each expert in that layer. Dashed lines (right) correspond to the average of layer 7 as depicted left. We display the first 32 experts out of 64. This plot is when $k = 1$ in Equation 8.

Vocabulary specialization refers to how specialized a particular expert is on a token ID x (also called a vocabulary element), defined as the proportion of tokens with a token ID x that are routed to one particular expert E_i out of all experts in that layer. We distinguish input and output variants of this specialization, where x is either the input token ID or the next output token ID (either the ground-truth next token ID or the token ID predicted by the model). A value of 100% indicates that for all occurrences of that vocabulary element, input data is routed to E_i , whereas 0% indicates an expert that is fully irrelevant for that vocabulary element and can be effectively removed from the model without affecting performance whenever the token ID appears.

In Figure 5 we find that vocabulary specialization is higher in later layers, similar to how later layers saturate earlier (§4.1). Later layers also specialize more on predicted output token IDs rather than input token IDs, i.e., the routing is decided more by the token the model is about to predict rather than

Table 4: **Vocabulary specialization in the 7th layer of OLMoE-1B-7B.** We use $k = 1$ (Equation 8) and a random 0.5% of the C4 validation data excluding token IDs with <10 appearances.

Expert ID	Input token IDs	Predicted output token IDs
27	© (100%) Ъ (100%) ³ (100%) í (100%) ĩ (100%) ̇ (100%) ̈́ (100%) क (100%) Ш (100%) ИH (100%) а (100%)	ı (100%) § (100%) © (100%) Ĵ (100%) ij ○ (100%) ① (100%) ₹ (100%) क (100%) ̇ (100%) ल (100%)
43	Armenian (100%) ijan (100%) enia (96%) Iraq (95%) Iranian (92%) Iran (92%) Saudi (90%) northern (90%) Lebanon (63%) (90%) Singapore (88%) Turkey (88%)	enia (90%) invasion (80%) Arabia (76%) irregular (66%) regions (64%) border Kong (61%) ians (61%) bases (60%) Republic (59%) Ireland (58%)
4	sq (89%) Main (70%) reversal (69%) YR (63%) GC (56%) Overall (50%) 79 (50%) main (50%) RE (46%) PCR (46%)	YR (90%) Character (88%) sq (77%) Os GHz (71%) fluence (60%) amycin (60%) pixels (56%) = (53%)
48	compared (42%) ! (41%) Then (41%) ', (40%)), (35%) ", (35%) instead (33%)	except (60%) tennis (41%) Marks (40%) Dunn (33%) tears (30%) Arizona (30%)
23 (58%) Therefore (55%) So (46%) !!! (46%) And (44%) According (41%) ." (47%) (41%) !! (40%) ?" (38%) But (38%)	Đ (53%) Republican (50%) Jack THIS (40%) Democratic (40%) according (39%) So (38%) Step (33%)

the original input token. This is intuitive as in earlier layers there is more uncertainty about which token the model will predict. At ~90%, expert 27 specializes the most, which we find in Table 4 to activate for many non-alphabetic tokens, such as Cyrillic and Devanagari letters. Expert 43 shows specialization on geographic terms in both input and output tokens. Experts 48 and 23 both focus on connector words, such as `Then` and `Therefore`. This is likely because they commonly process tokens together with a high co-activation of 60% in Figure 3. Based on our findings in §4.3 that for GitHub and arXiv often the same experts in layer 7 activate, we display one such expert (expert ID 4) in Table 4. It seems to specialize in measurements, such as `sq`, `YR` (year), and `GHz`. These are common terms in scientific papers corresponding to the arXiv domain and likely also in GitHub code for computations related to measurements. They are less likely to appear in books, which explains the low activation of expert ID 4 in layer 7 for book data in Figure 4. Expert 3 is among the three most active experts of layer 7 for book data in Figure 4 (fourth yellow bar for layer 7). This resonates when looking at its specialization on family terms in Table 4, which are far more common in books than scientific papers or code. Overall, domain specialization and vocabulary specialization are closely linked to one another, as domains are usually characterized by their distinct word distribution. In Figure H4, we link them more closely by comparing the extent of vocabulary specialization across domains and expert IDs. In Figure H2 and Figure H3 we also find that **OLMoE-1B-7B** exhibits stronger vocabulary specialization than Mixtral-8x7B.

5 Conclusion

We open-source **OLMoE-1B-7B** and **OLMoE-1B-7B-INSTRUCT** including model, data, code, and logs. At 1B active and 7B total parameters, our models yield state-of-the-art performance among models with a similar amount of active parameters even outperforming larger models including DeepSeekMoE-16B and Llama2-13B-Chat. We share ~20 training experiments yielding novel insights into Mixture-of-Experts. Further, we define and analyze new properties of MoEs showing that **OLMoE-1B-7B** exhibits early router saturation, weak expert co-activation, and some evidence of domain and vocabulary specialization. We intend our fully open release to serve as a basis for more research into MoEs given their critical importance (§1). We are excited about more iterations of **OLMoE** to close the gap between frontier models and fully open models.

Author Contributions

Niklas Muennighoff proposed and led the project. He ran the pretraining experiments, pretrained the model, helped run adaptation and analysis, and wrote most of the paper.

Luca Soldaini created the pretraining dataset and advised on data and architecture.

Dirk Groeneveld advised on pretraining, especially stability and throughput improvements.

Kyle Lo helped with pretraining dataset creation, analyzed data experiments, and advised on data and framing, and helped edit the paper.

Jacob Morrison co-created the adaptation dataset, ran most adaptation experiments, and helped edit the paper.

Sewon Min analyzed router saturation, expert correlation, and vocabulary specialization, and helped frame and edit the paper.

Weijia Shi analyzed domain and vocabulary specialization, advised at various project stages, and helped edit the paper.

Pete Walsh advised on pretraining, especially stability and throughput improvements.

Oyvind Tafjord ran OLMES evaluations.

Nathan Lambert co-created the adaptation dataset, advised on adaptation, and helped edit the paper.

Yuling Gu ran OLMES evaluations and helped edit the paper.

Shane Arora uploaded the models, helped with code review and framework integration.

Akshita Bhagia supported stability investigations and helped with DCLM evaluations.

Dustin Schwenk supported stability investigations.

David Wadden ran DCLM evaluations and helped with Weights & Biases reports.

Alexander Wettig analyzed load balancing, routing, and domain specialization, and helped edit the paper.

Binyuan Hui advised on pretraining and helped with plotting and framework integration.

Tim Dettmers advised on analysis and inference experiments.

Douwe Kiela advised on framing.

Ali Farhadi advised on pretraining and framing.

Noah A. Smith advised on pretraining, and helped frame and edit the paper.

Pang Wei Koh advised on analysis, and helped frame and edit the paper.

Amanpreet Singh advised on pretraining, framing and helped edit the paper.

Hannaneh Hajishirzi was responsible for direction and advising of the overall effort and helped frame and edit the paper.

Acknowledgements

OLMOE would not be possible without the support of many individuals and institutions. We thank our teammates at the Allen Institute for AI, Contextual AI, and the University of Washington for their support, especially Aditya Kusupati, Ananya Harsh Jha, Caitlin Wittlif, Carissa Schoenick, Costa Huang, Crystal Nam, David Atkinson, Emma Strubell, Faeze Brahman, Hamish Ivison, Karel D’Oosterlinck, Matt Latzke, Ian Magnusson, Jack Merullo, Jay Chen, Jennifer Dumas, Jiacheng Liu, Johann Dahm, Luke Zettlemoyer, Michael Schmitz, Michael Wilson, Pradeep Dasigi, Sahil Verma, Sam Skjonsberg, Sophie Lebrecht, Stas Bekman, Taira Anderson, Valentina Pyatkin, Yanai Elazar, Yizhong Wang, and Yoganand Chandrasekhar. We also thank Armen Aghajanyan, Akshat Shrivastava, Colin Raffel, Haokun Liu, Ludwig Schmidt, and Shayne Longpre. PWK is supported by the Singapore National Research Foundation and the National AI Group in the Singapore Ministry of Digital Development and Innovation under the AI Visiting Professorship Programme (award number AIVP-2024-001).

References

- [1] Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Qin Cai, Martin Cai, Caio César Teodoro Mendes, Weizhu Chen, Vishrav Chaudhary, Dong Chen, Dongdong Chen, Yen-Chun Chen, Yi-Ling Chen, Parul Chopra, Xiyang Dai, Allie Del Giorno, Gustavo de Rosa, Matthew Dixon, Ronen Eldan, Victor Fragoso, Dan Iter, Mei Gao, Min Gao, Jianfeng Gao, Amit Garg, Abhishek Goswami, Suriya Gunasekar, Emman Haider, Junheng Hao, Russell J. Hewett, Jamie Huynh, Mojan Javaheripi, Xin Jin, Piero Kauffmann, Nikos Karampatziakis, Dongwoo Kim, Mahoud Khademi, Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuanzhi Li, Yunsheng Li, Chen Liang, Lars Liden, Ce Liu, Mengchen Liu, Weishung Liu, Eric Lin, Zeqi Lin, Chong Luo, Piyush Madan, Matt Mazzola, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norick, Barun Patra, Daniel Perez-Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Corby Rosset, Sambudha Roy, Olatunji Ruwase, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacrose, Shital Shah, Ning Shang, Hiteshi Sharma, Swadheen Shukla, Xia Song, Masahiro Tanaka, Andrea Tupini, Xin Wang, Lijuan Wang, Chunyu Wang, Yu Wang, Rachel Ward, Guanhua Wang, Philipp Witte, Haiping Wu, Michael Wyatt, Bin Xiao, Can Xu, Jiahang Xu, Weijian Xu, Sonali Yadav, Fan Yang, Jianwei Yang, Ziyi Yang, Yifan Yang, Donghan Yu, Lu Yuan, Chengruidong 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](#).
- [2] 01. AI, :, Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, Kaidong Yu, Peng Liu, Qiang Liu, Shawn Yue, Senbin Yang, Shiming Yang, Tao Yu, Wen Xie, Wenhao Huang, Xiaohui Hu, Xiaoyi Ren, Xinyao Niu, Pengcheng Nie, Yuchi Xu, Yudong Liu, Yue Wang, Yuxuan Cai, Zhenyu Gu, Zhiyuan Liu, and Zonghong Dai. 2024. [Yi: Open Foundation Models by 01.AI](#).
- [3] Joshua Ainslie, James Lee-Thorp, Michiel de Jong, Yury Zemlyanskiy, Federico Lebrón, and Sumit Sanghai. 2023. [GQA: Training Generalized Multi-Query Transformer Models from Multi-Head Checkpoints](#).
- [4] Alon Albalak, Yanai Elazar, Sang Michael Xie, Shayne Longpre, Nathan Lambert, Xinyi Wang, Niklas Muennighoff, Bairu Hou, Liangming Pan, Haewon Jeong, Colin Raffel, Shiyu Chang, Tatsunori Hashimoto, and William Yang Wang. 2024. [A Survey on Data Selection for Language Models](#).
- [5] Loubna Ben Allal, Raymond Li, Denis Kocetkov, Chenghao Mou, Christopher Akiki, Carlos Munoz Ferrandis, Niklas Muennighoff, Mayank Mishra, Alex Gu, Manan Dey, et al. 2023. [SantaCoder: don't reach for the stars!](#)
- [6] Loubna Ben Allal, Anton Lozhkov, Elie Bakouch, Leandro von Werra, and Thomas Wolf. 2024. [SmolLM - blazingly fast and remarkably powerful](#).
- [7] Zeyuan Allen-Zhu and Yuanzhi Li. 2024. [Physics of Language Models: Part 3.3, Knowledge Capacity Scaling Laws](#).
- [8] Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Mérouane Debbah, Étienne Goffinet, Daniel Hesslow, Julien Launay, Quentin Malartic, Daniele Mazzotta, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. [The Falcon Series of Open Language Models](#).
- [9] Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepey, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, YaGuang Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni, Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussalem, Zachary Nado, John Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Aurko Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R. So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang,

- Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, and Yonghui Wu. 2023. [PaLM 2 Technical Report](#).
- [10] Mikel Artetxe, Shruti Bhosale, Naman Goyal, Todor Mihaylov, Myle Ott, Sam Shleifer, Xi Victoria Lin, Jingfei Du, Srinivasan Iyer, Ramakanth Pasunuru, Giri Anantharaman, Xian Li, Shuohui Chen, Halil Akin, Mandeep Baines, Louis Martin, Xing Zhou, Punit Singh Koura, Brian O’Horo, Jeff Wang, Luke Zettlemoyer, Mona Diab, Zornitsa Kozareva, and Ves Stoyanov. 2022. [Efficient Large Scale Language Modeling with Mixtures of Experts](#).
- [11] Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen McAleer, Albert Q. Jiang, Jia Deng, Stella Biderman, and Sean Welleck. 2023. [Llemma: An Open Language Model For Mathematics](#).
- [12] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer Normalization](#).
- [13] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. [Qwen Technical Report](#).
- [14] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. [Qwen-VL: A Versatile Vision-Language Model for Understanding, Localization, Text Reading, and Beyond](#).
- [15] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022. [Constitutional AI: Harmlessness from AI Feedback](#).
- [16] Marco Bellagente, Jonathan Tow, Dakota Mahan, Duy Phung, Maksym Zhuravinskiy, Reshinth Adithyan, James Baicoianu, Ben Brooks, Nathan Cooper, Ashish Datta, Meng Lee, Emad Mostaque, Michael Pieler, Nikhil Pinnaraju, Paulo Rocha, Harry Saini, Hannah Teufel, Niccolo Zanichelli, and Carlos Riquelme. 2024. [Stable LM 2 1.6B Technical Report](#).
- [17] Emmanuel Bengio, Pierre-Luc Bacon, Joelle Pineau, and Doina Precup. 2016. [Conditional Computation in Neural Networks for faster models](#).
- [18] Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar van der Wal. 2023. [Pythia: A Suite for Analyzing Large Language Models Across Training and Scaling](#).
- [19] Stella Biderman, Hailey Schoelkopf, Lintang Sutawika, Leo Gao, Jonathan Tow, Baber Abbasi, Alham Fikri Aji, Pawan Sasanka Ammanamanchi, Sidney Black, Jordan Clive, Anthony DiPofi, Julen Etxaniz, Benjamin Fattori, Jessica Zosa Forde, Charles Foster, Jeffrey Hsu, Mimansa Jaiswal, Wilson Y. Lee, Haonan Li, Charles Lovering, Niklas Muennighoff, Ellie Pavlick, Jason Phang, Aviya Skowron, Samson Tan, Xiangru Tang, Kevin A. Wang, Genta Indra Winata, François Yvon, and Andy Zou. 2024. [Lessons from the Trenches on Reproducible Evaluation of Language Models](#).
- [20] Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. 2019. [PIQA: Reasoning about Physical Commonsense in Natural Language](#).
- [21] Sid Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, Michael Pieler, USVSN Sai Prashanth, Shivanshu Purohit, Laria Reynolds, Jonathan Tow, Ben Wang, and Samuel Weinbach. 2022. [GPT-NeoX-20B: An Open-Source Autoregressive Language Model](#).
- [22] Sid Black, Leo Gao, Phil Wang, Connor Leahy, and Stella Biderman. 2021. [GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow](#).

- [23] Rishi Bommasani, Kevin Klyman, Shayne Longpre, Sayash Kapoor, Nestor Maslej, Betty Xiong, Daniel Zhang, and Percy Liang. 2023. [The Foundation Model Transparency Index](#).
- [24] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language Models are Few-Shot Learners](#).
- [25] Tianle Cai. 2023. [Mixtral from Mistral](#).
- [26] Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, Xiaoyi Dong, Haodong Duan, Qi Fan, Zhaoye Fei, Yang Gao, Jiaye Ge, Chenya Gu, Yuzhe Gu, Tao Gui, Aijia Guo, Qipeng Guo, Conghui He, Yingfan Hu, Ting Huang, Tao Jiang, Penglong Jiao, Zhenjiang Jin, Zhikai Lei, Jiaxing Li, Jingwen Li, Linyang Li, Shuaibin Li, Wei Li, Yining Li, Hongwei Liu, Jiangning Liu, Jiawei Hong, Kaiwen Liu, Kuikun Liu, Xiaoran Liu, Chengqi Lv, Haijun Lv, Kai Lv, Li Ma, Runyuan Ma, Zerun Ma, Wenchang Ning, Linke Ouyang, Jiantao Qiu, Yuan Qu, Fukai Shang, Yunfan Shao, Demin Song, Zifan Song, Zhihao Sui, Peng Sun, Yu Sun, Huanze Tang, Bin Wang, Guoteng Wang, Jiaqi Wang, Jiayu Wang, Rui Wang, Yudong Wang, Ziyi Wang, Xingjian Wei, Qizhen Weng, Fan Wu, Yingtong Xiong, Chao Xu, Ruiliang Xu, Hang Yan, Yirong Yan, Xiaogui Yang, Haochen Ye, Huaiyuan Ying, Jia Yu, Jing Yu, Yuhang Zang, Chuyu Zhang, Li Zhang, Pan Zhang, Peng Zhang, Ruijie Zhang, Shuo Zhang, Songyang Zhang, Wenjian Zhang, Wenwei Zhang, Xingcheng Zhang, Xinyue Zhang, Hui Zhao, Qian Zhao, Xiaomeng Zhao, Fengzhe Zhou, Zaida Zhou, Jingming Zhuo, Yicheng Zou, Xipeng Qiu, Yu Qiao, and Dahua Lin. 2024. [InternLM2 Technical Report](#).
- [27] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. 2020. [Generative pretraining from pixels](#).
- [28] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. [Evaluating Large Language Models Trained on Code](#).
- [29] Soumith Chintala. 2024. [GPT-4 MoE](#).
- [30] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. [PaLM: Scaling Language Modeling with Pathways](#).
- [31] Paul Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2023. [Deep reinforcement learning from human preferences](#).
- [32] Aidan Clark, Diego de las Casas, Aurelia Guy, Arthur Mensch, Michela Paganini, Jordan Hoffmann, Bogdan Damoc, Blake Hechtman, Trevor Cai, Sebastian Borgeaud, George van den Driessche, Eliza Rutherford, Tom Hennigan, Matthew Johnson, Katie Millican, Albin Cassirer, Chris Jones, Elena Buchatskaya, David Budden, Laurent Sifre, Simon Osindero, Oriol Vinyals, Jack Rae, Erich Elsen, Koray Kavukcuoglu, and Karen Simonyan. 2022. [Unified Scaling Laws for Routed Language Models](#).
- [33] 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](#).
- [34] Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have Solved Question Answering? Try ARC, the AI2 Reasoning Challenge](#).
- [35] 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. 2021. [Training Verifiers to Solve Math Word Problems](#).
- [36] Together Computer. 2023. [RedPajama: An Open Source Recipe to Reproduce LLaMA training dataset](#).
- [37] Róbert Csordás, Kazuki Irie, Jürgen Schmidhuber, Christopher Potts, and Christopher D. Manning. 2024. [MoEUT: Mixture-of-Experts Universal Transformers](#).
- [38] Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Wei Zhu, Yuan Ni, Guotong Xie, Zhiyuan Liu, and Maosong Sun. 2023. [UltraFeedback: Boosting Language Models with High-quality Feedback](#).

- [39] 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](#).
- [40] Databricks. 2024. [DBRX](#).
- [41] Yann N. Dauphin, Angela Fan, Michael Auli, and David Grangier. 2017. [Language Modeling with Gated Convolutional Networks](#).
- [42] DeepSeek-AI, :, Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiushi Du, Zhe Fu, Huazuo Gao, Kaige Gao, Wenjun Gao, Ruiqi Ge, Kang Guan, Daya Guo, Jianzhong Guo, Guangbo Hao, Zhewen Hao, Ying He, Wenjie Hu, Panpan Huang, Erhang Li, Guowei Li, Jiashi Li, Yao Li, Y. K. Li, Wenfeng Liang, Fangyun Lin, A. X. Liu, Bo Liu, Wen Liu, Xiaodong Liu, Xin Liu, Yiyuan Liu, Haoyu Lu, Shanghao Lu, Fuli Luo, Shirong Ma, Xiaotao Nie, Tian Pei, Yishi Piao, Junjie Qiu, Hui Qu, Tongzheng Ren, Zehui Ren, Chong Ruan, Zhangli Sha, Zhihong Shao, Junxiao Song, Xuecheng Su, Jingxiang Sun, Yaofeng Sun, Minghui Tang, Bingxuan Wang, Peiyi Wang, Shiyu Wang, Yaohui Wang, Yongji Wang, Tong Wu, Y. Wu, Xin Xie, Zhenda Xie, Ziwei Xie, Yiliang Xiong, Hanwei Xu, R. X. Xu, Yanhong Xu, Dejian Yang, Yuxiang You, Shuiping Yu, Xingkai Yu, B. Zhang, Haowei Zhang, Lecong Zhang, Liyue Zhang, Mingchuan Zhang, Minghua Zhang, Wentao Zhang, Yichao Zhang, Chenggang Zhao, Yao Zhao, Shangyan Zhou, Shunfeng Zhou, Qihao Zhu, and Yuheng Zou. 2024. [DeepSeek LLM: Scaling Open-Source Language Models with Longtermism](#).
- [43] DeepSeek-AI, Aixin Liu, Bei Feng, Bin Wang, Bingxuan Wang, Bo Liu, Chenggang Zhao, Chengqi Deng, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Hanwei Xu, Hao Yang, Haowei Zhang, Honghui Ding, Huaqian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi Ni, Jiashi Li, Jin Chen, Jingyang Yuan, Junjie Qiu, Junxiao Song, Kai Dong, Kaige Gao, Kang Guan, Lean Wang, Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Liyue Zhang, Meng Li, Miaojun Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan Huang, Peiyi Wang, Peng Zhang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J. Chen, R. L. Jin, Ruiqi Ge, Ruizhe Pan, Runxin Xu, Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shaoqing Wu, Shengfeng Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Size Zheng, T. Wang, Tian Pei, Tian Yuan, Tianyu Sun, W. L. Xiao, Wangding Zeng, Wei An, Wen Liu, Wenfeng Liang, Wenjun Gao, Wentao Zhang, X. Q. Li, Xiangyue Jin, Xianzu Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaojin Shen, Xiaokang Chen, Xiaosha Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Liu, Xin Xie, Xingkai Yu, Xinnan Song, Xinyi Zhou, Xinyu Yang, Xuan Lu, Xuecheng Su, Y. Wu, Y. K. Li, Y. X. Wei, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Li, Yaohui Wang, Yi Zheng, Yichao Zhang, Yiliang Xiong, Yilong Zhao, Ying He, Ying Tang, Yishi Piao, Yixin Dong, Yixuan Tan, Yiyuan Liu, Yongji Wang, Yongqiang Guo, Yuchen Zhu, Yudian Wang, Yuheng Zou, Yukun Zha, Yunxian Ma, Yuting Yan, Yuxiang You, Yuxuan Liu, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhen Huang, Zhen Zhang, Zhenda Xie, Zhewen Hao, Zhihong Shao, Zhiniu Wen, Zhipeng Xu, Zhongyu Zhang, Zhuoshu Li, Zihan Wang, Zihui Gu, Zilin Li, and Ziwei Xie. 2024. [DeepSeek-V2: A Strong, Economical, and Efficient Mixture-of-Experts Language Model](#).
- [44] Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin Gilmer, Andreas Steiner, Mathilde Caron, Robert Geirhos, Ibrahim Alabdulmohsin, Rodolphe Jenatton, Lucas Beyer, Michael Tschannen, Anurag Arnab, Xiao Wang, Carlos Riquelme, Matthias Minderer, Joan Puigcerver, Utku Evci, Manoj Kumar, Sjoerd van Steenkiste, Gamaleldin F. Elsayed, Aravindh Mahendran, Fisher Yu, Avital Oliver, Fantine Huot, Jasmijn Bastings, Mark Patrick Collier, Alexey Gritsenko, Vighnesh Birodkar, Cristina Vasconcelos, Yi Tay, Thomas Mensink, Alexander Kolesnikov, Filip Pavetić, Dustin Tran, Thomas Kipf, Mario Lučić, Xiaohua Zhai, Daniel Keysers, Jeremiah Harmsen, and Neil Houlsby. 2023. [Scaling Vision Transformers to 22 Billion Parameters](#).
- [45] Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. 2019. [Universal Transformers](#).
- [46] Nolan Dey, Gurpreet Gosal, Zhiming Chen, Hemant Khachane, William Marshall, Ribhu Pathria, Marvin Tom, and Joel Hestness. 2023. [Cerebras-GPT: Open Compute-Optimal Language Models Trained on the Cerebras Wafer-Scale Cluster](#).
- [47] Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Pete Florence. 2023. [PaLM-E: An Embodied Multimodal Language Model](#).

- [48] Nan Du, Yanping Huang, Andrew M. Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, Barret Zoph, Liam Fedus, Maarten Bosma, Zongwei Zhou, Tao Wang, Yu Emma Wang, Kellie Webster, Marie Pellat, Kevin Robinson, Kathleen Meier-Hellstern, Toju Duke, Lucas Dixon, Kun Zhang, Quoc V Le, Yonghui Wu, Zhifeng Chen, and Claire Cui. 2022. [GLaM: Efficient Scaling of Language Models with Mixture-of-Experts](#).
- [49] Dheeru Dua, Shruti Bhosale, Vedanuj Goswami, James Cross, Mike Lewis, and Angela Fan. 2021. [Tricks for Training Sparse Translation Models](#).
- [50] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, et al. 2024. [The Llama 3 Herd of Models](#).
- [51] Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B. Hashimoto. 2024. [Length-Controlled AlpacaEval: A Simple Way to Debias Automatic Evaluators](#).
- [52] David Eigen, Marc’ Aurelio Ranzato, and Ilya Sutskever. 2014. [Learning Factored Representations in a Deep Mixture of Experts](#).
- [53] Kenneth Enevoldsen, Márton Kardos, Niklas Muennighoff, and Kristoffer Laigaard Nielbo. 2024. [The Scandinavian Embedding Benchmarks: Comprehensive Assessment of Multilingual and Monolingual Text Embedding](#).
- [54] Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. [KTO: Model Alignment as Prospect Theoretic Optimization](#).
- [55] Manuel Faysse, Patrick Fernandes, Nuno M. Guerreiro, António Loison, Duarte M. Alves, Caio Corro, Nicolas Boizard, João Alves, Ricardo Rei, Pedro H. Martins, Antoni Bigata Casademunt, François Yvon, André F. T. Martins, Gautier Viaud, Céline Hudelot, and Pierre Colombo. 2024. [CroissantLLM: A Truly Bilingual French-English Language Model](#).
- [56] William Fedus, Barret Zoph, and Noam Shazeer. 2022. [Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity](#).
- [57] Samir Yitzhak Gadre, Georgios Smyrnis, Vaishaal Shankar, Suchin Gururangan, Mitchell Wortsman, Rulin Shao, Jean Mercat, Alex Fang, Jeffrey Li, Sedrick Keh, Rui Xin, Marianna Nezhurina, Igor Vasiljevic, Jenia Jitsev, Luca Soldaini, Alexandros G. Dimakis, Gabriel Ilharco, Pang Wei Koh, Shuran Song, Thomas Kollar, Yair Carmon, Achal Dave, Reinhard Heckel, Niklas Muennighoff, and Ludwig Schmidt. 2024. [Language models scale reliably with over-training and on downstream tasks](#).
- [58] Trevor Gale, Deepak Narayanan, Cliff Young, and Matei Zaharia. 2022. [MegaBlocks: Efficient Sparse Training with Mixture-of-Experts](#).
- [59] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2020. [The Pile: An 800GB Dataset of Diverse Text for Language Modeling](#).
- [60] Leo Gao, Jonathan Tow, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Kyle McDonell, Niklas Muennighoff, Jason Phang, Laria Reynolds, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2021. [A framework for few-shot language model evaluation](#).
- [61] Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiada Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, Juanzi Li, Lei Zhao, Lindong Wu, Lucen Zhong, Mingdao Liu, Minlie Huang, Peng Zhang, Qinkai Zheng, Rui Lu, Shuaiqi Duan, Shudan Zhang, Shulin Cao, Shuxun Yang, Weng Lam Tam, Wenyi Zhao, Xiao Liu, Xiao Xia, Xiaohan Zhang, Xiaotao Gu, Xin Lv, Xinghan Liu, Xinyi Liu, Xinyue Yang, Xixuan Song, Xunkai Zhang, Yifan An, Yifan Xu, Yilin Niu, Yuantao Yang, Yueyan Li, Yushi Bai, Yuxiao Dong, Zehan Qi, Zhaoyu Wang, Zhen Yang, Zhengxiao Du, Zhenyu Hou, and Zihan Wang. 2024. [ChatGLM: A Family of Large Language Models from GLM-130B to GLM-4 All Tools](#).
- [62] Paolo Glorioso, Quentin Anthony, Yury Tokpanov, James Whittington, Jonathan Pilault, Adam Ibrahim, and Beren Millidge. 2024. [Zamba: A Compact 7B SSM Hybrid Model](#).
- [63] Andrew Gordon, Zornitsa Kozareva, and Melissa Roemmele. 2012. [SemEval-2012 Task 7: Choice of Plausible Alternatives: An Evaluation of Commonsense Causal Reasoning](#).

- [64] Dirk Groeneveld, Anas Awadalla, Iz Beltagy, Akshita Bhagia, Ian Magnusson, Hao Peng, Oyvind Tafjord, Pete Walsh, Kyle Richardson, and Jesse Dodge. 2023. [Catwalk: A Unified Language Model Evaluation Framework for Many Datasets](#).
- [65] Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Raghavi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, Will Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah A. Smith, and Hannaneh Hajishirzi. 2024. [OLMo: Accelerating the Science of Language Models](#).
- [66] Sam Gross, Marc’Aurelio Ranzato, and Arthur Szlam. 2017. [Hard Mixtures of Experts for Large Scale Weakly Supervised Vision](#).
- [67] Yuling Gu, Oyvind Tafjord, Bailey Kuehl, Dany Haddad, Jesse Dodge, and Hannaneh Hajishirzi. 2024. [OLMES: A Standard for Language Model Evaluations](#).
- [68] Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. 2023. [Textbooks Are All You Need](#).
- [69] Xu Owen He. 2024. [Mixture of A Million Experts](#).
- [70] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring Massive Multitask Language Understanding](#).
- [71] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring Mathematical Problem Solving With the MATH Dataset](#).
- [72] 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. 2022. [Training Compute-Optimal Large Language Models](#).
- [73] Jiwoo Hong, Noah Lee, and James Thorne. 2024. [ORPO: Monolithic Preference Optimization without Reference Model](#).
- [74] Shengding Hu, Yuge Tu, Xu Han, Chaoqun He, Ganqu Cui, Xiang Long, Zhi Zheng, Yewei Fang, Yuxiang Huang, Weilin Zhao, Xinrong Zhang, Zheng Leng Thai, Kaihuo Zhang, Chongyi Wang, Yuan Yao, Chenyang Zhao, Jie Zhou, Jie Cai, Zhongwu Zhai, Ning Ding, Chao Jia, Guoyang Zeng, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024. [MiniCPM: Unveiling the Potential of Small Language Models with Scalable Training Strategies](#).
- [75] Cheng-Zhi Anna Huang, Ashish Vaswani, Jakob Uszkoreit, Noam Shazeer, Ian Simon, Curtis Hawthorne, Andrew M. Dai, Matthew D. Hoffman, Monica Dinulescu, and Douglas Eck. 2018. [Music Transformer](#).
- [76] Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. 2023. [Camels in a Changing Climate: Enhancing LM Adaptation with Tulu 2](#).
- [77] Sebastian Jaszczur, Aakanksha Chowdhery, Afroz Mohiuddin, Łukasz Kaiser, Wojciech Gajewski, Henryk Michalewski, and Jonni Kanerva. 2021. [Sparse is Enough in Scaling Transformers](#).
- [78] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7B](#).
- [79] 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. 2024. [Mixtral of Experts](#).

- [80] 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](#).
- [81] Andrej Karpathy. 2024. [LLM model size competition is intensifying. . . backwards!](#)
- [82] Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Casey A Fitzpatrick, Peter Bull, Greg Lipstein, Tony Nelli, Ron Zhu, et al. 2021. [The hateful memes challenge: Competition report](#).
- [83] Diederik P. Kingma and Jimmy Ba. 2017. [Adam: A Method for Stochastic Optimization](#).
- [84] Denis Kocetkov, Raymond Li, Loubna Ben Allal, Jia Li, Chenghao Mou, Carlos Muñoz Ferrandis, Yacine Jernite, Margaret Mitchell, Sean Hughes, Thomas Wolf, Dzmitry Bahdanau, Leandro von Werra, and Harm de Vries. 2022. [The Stack: 3 TB of permissively licensed source code](#).
- [85] Aran Komatsuzaki, Joan Puigcerver, James Lee-Thorp, Carlos Riquelme Ruiz, Basil Mustafa, Joshua Ainslie, Yi Tay, Mostafa Dehghani, and Neil Houlsby. 2023. [Sparse Upcycling: Training Mixture-of-Experts from Dense Checkpoints](#).
- [86] Jakub Krajewski, Jan Ludziejewski, Kamil Adamczewski, Maciej Pióro, Michał Krutul, Szymon Antoniak, Kamil Ciebiera, Krystian Król, Tomasz Odrzygóźdź, Piotr Sankowski, Marek Cygan, and Sebastian Jaszczur. 2024. [Scaling Laws for Fine-Grained Mixture of Experts](#).
- [87] 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](#).
- [88] Mike Lewis, Shruti Bhosale, Tim Dettmers, Naman Goyal, and Luke Zettlemoyer. 2021. [BASE Layers: Simplifying Training of Large, Sparse Models](#).
- [89] Jeffrey Li, Alex Fang, Georgios Smyrnis, Maor Ivgi, Matt Jordan, Samir Gadre, Hritik Bansal, Etash Guha, Sedrick Keh, Kushal Arora, Saurabh Garg, Rui Xin, Niklas Muennighoff, Reinhard Heckel, Jean Mercat, Mayee Chen, Suchin Gururangan, Mitchell Wortsman, Alon Albalak, Yonatan Bitton, Marianna Nezhurina, Amro Abbas, Cheng-Yu Hsieh, Dhruva Ghosh, Josh Gardner, Maciej Kilian, Hanlin Zhang, Rulin Shao, Sarah Pratt, Sunny Sanyal, Gabriel Ilharco, Giannis Daras, Kalyani Marathe, Aaron Gokaslan, Jieyu Zhang, Khyathi Chandu, Thao Nguyen, Igor Vasiljevic, Sham Kakade, Shuran Song, Sujay Sanghavi, Fartash Faghri, Sewoong Oh, Luke Zettlemoyer, Kyle Lo, Alaaeldin El-Nouby, Hadi Pouransari, Alexander Toshev, Stephanie Wang, Dirk Groeneveld, Luca Soldaini, Pang Wei Koh, Jenia Jitsev, Thomas Kollar, Alexandros G. Dimakis, Yair Carmon, Achal Dave, Ludwig Schmidt, and Vaishaal Shankar. 2024. [DataComp-LM: In search of the next generation of training sets for language models](#).
- [90] Margaret Li, Suchin Gururangan, Tim Dettmers, Mike Lewis, Tim Althoff, Noah A. Smith, and Luke Zettlemoyer. 2022. [Branch-Train-Merge: Embarrassingly Parallel Training of Expert Language Models](#).
- [91] Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, et al. 2023. [StarCoder: may the source be with you!](#)
- [92] Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [AlpacaEval: An Automatic Evaluator of Instruction-following Models](#).
- [93] Yuanzhi Li, Sébastien Bubeck, Ronen Eldan, Allie Del Giorno, Suriya Gunasekar, and Yin Tat Lee. 2023. [Textbooks Are All You Need II: phi-1.5 technical report](#).
- [94] Yunxin Li, Shenyan Jiang, Baotian Hu, Longyue Wang, Wanqi Zhong, Wenhan Luo, Lin Ma, and Min Zhang. 2024. [Uni-MoE: Scaling Unified Multimodal LLMs with Mixture of Experts](#).
- [95] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue Wang, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekogul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2023. [Holistic Evaluation of Language Models](#).

- [96] Opher Lieber, Barak Lenz, Hofit Bata, Gal Cohen, Jhonathan Osin, Itay Dalmedigos, Erez Safahi, Shaked Meirom, Yonatan Belinkov, Shai Shalev-Shwartz, Omri Abend, Raz Alon, Tomer Asida, Amir Bergman, Roman Glozman, Michael Gokhman, Avashalom Manevich, Nir Ratner, Noam Rozen, Erez Shwartz, Mor Zusman, and Yoav Shoham. 2024. [Jamba: A Hybrid Transformer-Mamba Language Model](#).
- [97] Bin Lin, Zhenyu Tang, Yang Ye, Jiayi Cui, Bin Zhu, Peng Jin, Jinfa Huang, Junwu Zhang, Yatian Pang, Munan Ning, and Li Yuan. 2024. [MoE-LLaVA: Mixture of Experts for Large Vision-Language Models](#).
- [98] Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [TruthfulQA: Measuring How Models Mimic Human Falsehoods](#).
- [99] Xi Victoria Lin, Akshat Shrivastava, Liang Luo, Srinivasan Iyer, Mike Lewis, Gargi Gosh, Luke Zettlemoyer, and Armen Aghajanyan. 2024. [MoMa: Efficient Early-Fusion Pre-training with Mixture of Modality-Aware Experts](#).
- [100] Qian Liu, Xiaosen Zheng, Niklas Muennighoff, Guangtao Zeng, Longxu Dou, Tianyu Pang, Jing Jiang, and Min Lin. 2024. [RegMix: Data Mixture as Regression for Language Model Pre-training](#).
- [101] Tianlin Liu, Mathieu Blondel, Carlos Riquelme, and Joan Puigcerver. 2024. [Routes in Vision Mixture of Experts: An Empirical Study](#).
- [102] Zhengzhong Liu, Aurick Qiao, Willie Neiswanger, Hongyi Wang, Bowen Tan, Tianhua Tao, Junbo Li, Yuqi Wang, Suqi Sun, Omkar Pangarkar, Richard Fan, Yi Gu, Victor Miller, Yonghao Zhuang, Guowei He, Haonan Li, Fajri Koto, Liping Tang, Nikhil Ranjan, Zhiqiang Shen, Xuguang Ren, Roberto Iriondo, Cun Mu, Zhiting Hu, Mark Schulze, Preslav Nakov, Tim Baldwin, and Eric P. Xing. 2023. [LLM360: Towards Fully Transparent Open-Source LLMs](#).
- [103] Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V. Le, Barret Zoph, Jason Wei, and Adam Roberts. 2023. [The Flan Collection: Designing Data and Methods for Effective Instruction Tuning](#).
- [104] Shayne Longpre, Robert Mahari, Anthony Chen, Naana Obeng-Marnu, Damien Sileo, William Brannon, Niklas Muennighoff, Nathan Khazam, Jad Kabbara, Kartik Perisetla, Xinyi Wu, Enrico Shippole, Kurt Bollacker, Tongshuang Wu, Luis Villa, Sandy Pentland, and Sara Hooker. 2023. [The Data Provenance Initiative: A Large Scale Audit of Dataset Licensing & Attribution in AI](#).
- [105] Shayne Longpre, Robert Mahari, Ariel Lee, Campbell Lund, Hamidah Oderinwale, William Brannon, Nayan Saxena, Naana Obeng-Marnu, Tobin South, Cole Hunter, Kevin Klyman, Christopher Klamm, Hailey Schoelkopf, Nikhil Singh, Manuel Cherep, Ahmad Anis, An Dinh, Caroline Chitongo, Da Yin, Damien Sileo, Deividas Mataciunas, Diganta Misra, Emad Alghamdi, Enrico Shippole, Jianguo Zhang, Joanna Materzynska, Kun Qian, Kush Tiwary, Lester Miranda, Manan Dey, Minnie Liang, Mohammed Hamdy, Niklas Muennighoff, Seonghyeon Ye, Seungone Kim, Shrestha Mohanty, Vipul Gupta, Vivek Sharma, Vu Minh Chien, Xuhui Zhou, Yizhi Li, Caiming Xiong, Luis Villa, Stella Biderman, Hanlin Li, Daphne Ippolito, Sara Hooker, Jad Kabbara, and Sandy Pentland. 2024. [Consent in Crisis: The Rapid Decline of the AI Data Commons](#).
- [106] Ilya Loshchilov and Frank Hutter. 2019. [Decoupled Weight Decay Regularization](#).
- [107] Holy Lovenia, Rahmad Mahendra, Salsabil Maulana Akbar, Lester James V. Miranda, Jennifer Santoso, Elyanah Aco, Akhdan Fadhillah, Jonibek Mansurov, Joseph Marvin Imperial, Onno P. Kampman, Joel Ruben Antony Moniz, Muhammad Ravi Shulthan Habibi, Frederikus Hudi, Railey Montalan, Ryan Ignatius, Joanito Agili Lopo, William Nixon, Börje F. Karlsson, James Jaya, Ryandito Diandaru, Yuze Gao, Patrick Amadeus, Bin Wang, Jan Christian Blaise Cruz, Chenxi Whitehouse, Ivan Halim Parmonangan, Maria Khelli, Wenyu Zhang, Lucky Susanto, Reynard Adha Ryanda, Sonny Lazuardi Hermawan, Dan John Velasco, Muhammad Dehan Al Kautsar, Willy Fitra Hendria, Yasmin Moslem, Noah Flynn, Muhammad Farid Adilazuarda, Haochen Li, Johannes Lee, R. Damanhuri, Shuo Sun, Muhammad Reza Qorib, Amirbek Djanibekov, Wei Qi Leong, Quyet V. Do, Niklas Muennighoff, Tanrada Pansuwan, Ilham Firdausi Putra, Yan Xu, Ngee Chia Tai, Ayu Purwarianti, Sebastian Ruder, William Tjhi, Peerat Limkonchotiwat, Alham Fikri Aji, Sedrick Keh, Genta Indra Winata, Ruochen Zhang, Fajri Koto, Zheng-Xin Yong, and Samuel Cahyawijaya. 2024. [SEACrowd: A Multilingual Multimodal Data Hub and Benchmark Suite for Southeast Asian Languages](#).
- [108] Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, Tianyang Liu, Max Tian, Denis Kocetkov, Arthur Zucker, Younes Belkada, Zijian Wang, Qian Liu, Dmitry Abulkhanov, Indraneil Paul, Zhuang Li, Wen-Ding Li, Megan Risdal, Jia Li, Jian Zhu, Terry Yue Zhuo, Evgenii Zheltonozhskii, Nii Osae Osae Dade, Wenhao Yu, Lucas Krauß, Naman Jain, Yixuan Su, Xuanli He, Manan Dey, Edoardo Abati, Yekun

- Chai, Niklas Muennighoff, Xiangru Tang, Muhtasham Oblokulov, Christopher Akiki, Marc Marone, Chenghao Mou, Mayank Mishra, Alex Gu, Binyuan Hui, Tri Dao, Armel Zebaze, Olivier Dehaene, Nicolas Patry, Canwen Xu, Julian McAuley, Han Hu, Torsten Scholak, Sebastien Paquet, Jennifer Robinson, Carolyn Jane Anderson, Nicolas Chapados, Mostofa Patwary, Nima Tajbakhsh, Yacine Jernite, Carlos Muñoz Ferrandis, Lingming Zhang, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. 2024. [StarCoder 2 and The Stack v2: The Next Generation](#).
- [109] Risto Luukkonen, Ville Komulainen, Jouni Luoma, Anni Eskelinen, Jenna Kanerva, Hanna-Mari Kupari, Filip Ginter, Veronika Laippala, Niklas Muennighoff, Aleksandra Piktus, Thomas Wang, Nouamane Tazi, Teven Le Scao, Thomas Wolf, Osma Suominen, Samuli Sairanen, Mikko Merioksa, Jyrki Heinonen, Aija Vahtola, Samuel Antao, and Sampo Pyysalo. 2023. [FinGPT: Large Generative Models for a Small Language](#).
- [110] Ian Magnusson, Akshita Bhagia, Valentin Hofmann, Luca Soldaini, Ananya Harsh Jha, Oyvind Tafjord, Dustin Schwenk, Evan Pete Walsh, Yanai Elazar, Kyle Lo, Dirk Groeneveld, Iz Beltagy, Hannaneh Hajishirzi, Noah A. Smith, Kyle Richardson, and Jesse Dodge. 2023. [Paloma: A Benchmark for Evaluating Language Model Fit](#).
- [111] Brandon McKinzie, Zhe Gan, Jean-Philippe Fauconnier, Sam Dodge, Bowen Zhang, Philipp Dufer, Dhruvi Shah, Xianzhi Du, Futang Peng, Floris Weers, Anton Belyi, Haotian Zhang, Karanjeet Singh, Doug Kang, Ankur Jain, Hongyu He, Max Schwarzer, Tom Gunter, Xiang Kong, Aonan Zhang, Jianyu Wang, Chong Wang, Nan Du, Tao Lei, Sam Wiseman, Guoli Yin, Mark Lee, Zirui Wang, Ruoming Pang, Peter Gräsch, Alexander Toshev, and Yinfei Yang. 2024. [MM1: Methods, Analysis & Insights from Multimodal LLM Pre-training](#).
- [112] Sachin Mehta, Mohammad Hossein Sekhvat, Qingqing Cao, Maxwell Horton, Yanzi Jin, Chenfan Sun, Iman Mirzadeh, Mahyar Najibi, Dmitry Belenko, Peter Zatloukal, and Mohammad Rastegari. 2024. [OpenELM: An Efficient Language Model Family with Open Training and Inference Framework](#).
- [113] Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. [SimPO: Simple Preference Optimization with a Reference-Free Reward](#).
- [114] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. [Pointer Sentinel Mixture Models](#).
- [115] Paulius Micekevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. 2018. [Mixed Precision Training](#).
- [116] Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. [Can a Suit of Armor Conduct Electricity? A New Dataset for Open Book Question Answering](#).
- [117] Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2022. [Cross-Task Generalization via Natural Language Crowdsourcing Instructions](#).
- [118] Niklas Muennighoff. 2020. [Vilio: State-of-the-art Visio-Linguistic Models applied to Hateful Memes](#).
- [119] Niklas Muennighoff, Qian Liu, Armel Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo, Swayam Singh, Xiangru Tang, Leandro von Werra, and Shayne Longpre. 2023. [OctoPack: Instruction Tuning Code Large Language Models](#).
- [120] Niklas Muennighoff, Alexander M. Rush, Boaz Barak, Teven Le Scao, Aleksandra Piktus, Nouamane Tazi, Sampo Pyysalo, Thomas Wolf, and Colin Raffel. 2023. [Scaling Data-Constrained Language Models](#).
- [121] Niklas Muennighoff, Hongjin Su, Liang Wang, Nan Yang, Furu Wei, Tao Yu, Amanpreet Singh, and Douwe Kiela. 2024. [Generative Representational Instruction Tuning](#).
- [122] Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 2023. [Crosslingual Generalization through Multitask Finetuning](#).
- [123] Mohammed Muqeeth, Haokun Liu, and Colin Raffel. 2024. [Soft Merging of Experts with Adaptive Routing](#).
- [124] Basil Mustafa, Carlos Riquelme, Joan Puigcerver, Rodolphe Jenatton, and Neil Houlsby. 2022. [Multimodal Contrastive Learning with LIMoE: the Language-Image Mixture of Experts](#).

- [125] Nvidia, :, Bo Adler, Niket Agarwal, Ashwath Aithal, Dong H. Anh, Pallab Bhattacharya, Annika Brundyn, Jared Casper, Bryan Catanzaro, Sharon Clay, Jonathan Cohen, Sirshak Das, Ayush Dattagupta, Olivier Delalleau, Leon Derczynski, Yi Dong, Daniel Egert, Ellie Evans, Aleksander Ficek, Denys Fridman, Shaona Ghosh, Boris Ginsburg, Igor Gitman, Tomasz Grzegorzec, Robert Hero, Jining Huang, Vibhu Jawa, Joseph Jennings, Aastha Jhunjhunwala, John Kamalu, Sadaf Khan, Oleksii Kuchaiev, Patrick LeGresley, Hui Li, Jiwei Liu, Zihan Liu, Eileen Long, Ameya Sunil Mahabaleshwarkar, Somshubra Majumdar, James Maki, Miguel Martinez, Maer Rodrigues de Melo, Ivan Moshkov, Deepak Narayanan, Sean Narenthiran, Jesus Navarro, Phong Nguyen, Osvald Nitski, Vahid Noroozi, Guruprasad Nutheti, Christopher Parisien, Jupinder Parmar, Mostofa Patwary, Krzysztof Pawelec, Wei Ping, Shrimai Prabhummoye, Rajarshi Roy, Trisha Saar, Vasanth Rao Naik Sabavat, Sanjeev Satheesh, Jane Polak Scowcroft, Jason Sewall, Pavel Shamis, Gerald Shen, Mohammad Shoeybi, Dave Sizer, Misha Smelyanskiy, Felipe Soares, Makesh Narsimhan Sreedhar, Dan Su, Sandeep Subramanian, Shengyang Sun, Shubham Toshniwal, Hao Wang, Zhilin Wang, Jiakuan You, Jiaqi Zeng, Jimmy Zhang, Jing Zhang, Vivienne Zhang, Yian Zhang, and Chen Zhu. 2024. [Nemotron-4 340B Technical Report](#).
- [126] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, et al. 2023. [GPT-4 Technical Report](#).
- [127] Bowen Pan, Yikang Shen, Haokun Liu, Mayank Mishra, Gaoyuan Zhang, Aude Oliva, Colin Raffel, and Rameswar Panda. 2024. [Dense Training, Sparse Inference: Rethinking Training of Mixture-of-Experts Language Models](#).
- [128] Jupinder Parmar, Shrimai Prabhummoye, Joseph Jennings, Mostofa Patwary, Sandeep Subramanian, Dan Su, Chen Zhu, Deepak Narayanan, Aastha Jhunjhunwala, Ayush Dattagupta, Vibhu Jawa, Jiwei Liu, Ameya Mahabaleshwarkar, Osvald Nitski, Annika Brundyn, James Maki, Miguel Martinez, Jiakuan You, John Kamalu, Patrick LeGresley, Denys Fridman, Jared Casper, Ashwath Aithal, Oleksii Kuchaiev, Mohammad Shoeybi, Jonathan Cohen, and Bryan Catanzaro. 2024. [Nemotron-4 15B Technical Report](#).
- [129] Keiran Paster, Marco Dos Santos, Zhangir Azerbayev, and Jimmy Ba. 2023. [OpenWebMath: An Open Dataset of High-Quality Mathematical Web Text](#).
- [130] Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtessam Almazrouei, and Julien Launay. 2023. [The RefinedWeb Dataset for Falcon LLM: Outperforming Curated Corpora with Web Data, and Web Data Only](#).
- [131] Bo Peng, Eric Alcaide, Quentin Anthony, Alon Albalak, Samuel Arcadinho, Stella Biderman, Huanqi Cao, Xin Cheng, Michael Chung, Matteo Grella, Kranthi Kiran GV, Xuzheng He, Haowen Hou, Jiaju Lin, Przemyslaw Kazienko, Jan Kocon, Jiaming Kong, Bartlomiej Koptyra, Hayden Lau, Krishna Sri Ipsit Mantri, Ferdinand Mom, Atsushi Saito, Guangyu Song, Xiangru Tang, Bolun Wang, Johan S. Wind, Stanislaw Wozniak, Ruichong Zhang, Zhenyuan Zhang, Qihang Zhao, Peng Zhou, Qinghua Zhou, Jian Zhu, and Rui-Jie Zhu. 2023. [RWKV: Reinventing RNNs for the Transformer Era](#).
- [132] Bo Peng, Daniel Goldstein, Quentin Anthony, Alon Albalak, Eric Alcaide, Stella Biderman, Eugene Cheah, Xingjian Du, Teddy Ferdinan, Haowen Hou, Przemyslaw Kazienko, Kranthi Kiran GV, Jan Kocoń, Bartlomiej Koptyra, Satyapriya Krishna, Ronald McClelland Jr. au2, Niklas Muennighoff, Fares Obeid, Atsushi Saito, Guangyu Song, Haoqin Tu, Stanislaw Wozniak, Ruichong Zhang, Bingchen Zhao, Qihang Zhao, Peng Zhou, Jian Zhu, and Rui-Jie Zhu. 2024. [Eagle and Finch: RWKV with Matrix-Valued States and Dynamic Recurrence](#).
- [133] Ofir Press and Lior Wolf. 2017. [Using the Output Embedding to Improve Language Models](#).
- [134] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. [Robust Speech Recognition via Large-Scale Weak Supervision](#).
- [135] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. [Language models are unsupervised multitask learners](#).
- [136] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2023. [Direct Preference Optimization: Your Language Model is Secretly a Reward Model](#).
- [137] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2023. [Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer](#).
- [138] Nazneen Rajani, Lewis Tunstall, Edward Beeching, Nathan Lambert, Alexander M. Rush, and Thomas Wolf. 2023. [No Robots](#).

- [139] Samyam Rajbhandari, Conglong Li, Zhewei Yao, Minjia Zhang, Reza Yazdani Aminabadi, Ammar Ahmad Awan, Jeff Rasley, and Yuxiong He. 2022. [DeepSpeed-MoE: Advancing Mixture-of-Experts Inference and Training to Power Next-Generation AI Scale.](#)
- [140] Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. [ZeRO: Memory Optimizations Toward Training Trillion Parameter Models.](#)
- [141] David Raposo, Sam Ritter, Blake Richards, Timothy Lillicrap, Peter Conway Humphreys, and Adam Santoro. 2024. [Mixture-of-Depths: Dynamically allocating compute in transformer-based language models.](#)
- [142] Machel Reid, Victor Zhong, Suchin Gururangan, and Luke Zettlemoyer. 2022. [M2D2: A Massively Multi-domain Language Modeling Dataset.](#)
- [143] Xiaozhe Ren, Pingyi Zhou, Xinfan Meng, Xinjing Huang, Yadao Wang, Weichao Wang, Pengfei Li, Xiaoda Zhang, Alexander Podolskiy, Grigory Arshinov, Andrey Bout, Irina Piontkovskaya, Jiansheng Wei, Xin Jiang, Teng Su, Qun Liu, and Jun Yao. 2023. [PanGu-Sigma: Towards Trillion Parameter Language Model with Sparse Heterogeneous Computing.](#)
- [144] Stephen Roller, Sainbayar Sukhbaatar, Arthur Szlam, and Jason Weston. 2021. [Hash Layers For Large Sparse Models.](#)
- [145] Paul Röttger, Hannah Rose Kirk, Bertie Vidgen, Giuseppe Attanasio, Federico Bianchi, and Dirk Hovy. 2024. [XSTest: A Test Suite for Identifying Exaggerated Safety Behaviours in Large Language Models.](#)
- [146] Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. [WinoGrande: An Adversarial Winograd Schema Challenge at Scale.](#)
- [147] Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, et al. 2022. [Multitask Prompted Training Enables Zero-Shot Task Generalization.](#)
- [148] Maarten Sap, Hannah Rashkin, Derek Chen, Ronan LeBras, and Yejin Choi. 2019. [SocialQA: Commonsense Reasoning about Social Interactions.](#)
- [149] Teven Le Scao, Thomas Wang, Daniel Hesslow, Lucile Saulnier, Stas Bekman, M Saiful Bari, Stella Biderman, Hady Elsahar, Niklas Muennighoff, Jason Phang, Ofir Press, Colin Raffel, Victor Sanh, Sheng Shen, Lintang Sutawika, Jaesung Tae, Zheng Xin Yong, Julien Launay, and Iz Beltagy. 2022. [What Language Model to Train if You Have One Million GPU Hours?](#)
- [150] Noam Shazeer. 2019. [Fast Transformer Decoding: One Write-Head is All You Need.](#)
- [151] Noam Shazeer. 2020. [GLU Variants Improve Transformer.](#)
- [152] 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.](#)
- [153] Noam Shazeer and Mitchell Stern. 2018. [Adafactor: Adaptive Learning Rates with Sublinear Memory Cost.](#)
- [154] Sheng Shen, Le Hou, Yanqi Zhou, Nan Du, Shayne Longpre, Jason Wei, Hyung Won Chung, Barret Zoph, William Fedus, Xinyun Chen, Tu Vu, Yuexin Wu, Wuyang Chen, Albert Webson, Yunxuan Li, Vincent Zhao, Hongkun Yu, Kurt Keutzer, Trevor Darrell, and Denny Zhou. 2023. [Mixture-of-Experts Meets Instruction Tuning: A Winning Combination for Large Language Models.](#)
- [155] Sheng Shen, Zhewei Yao, Chunyuan Li, Trevor Darrell, Kurt Keutzer, and Yuxiong He. 2023. [Scaling Vision-Language Models with Sparse Mixture of Experts.](#)
- [156] Yikang Shen, Zhen Guo, Tianle Cai, and Zengyi Qin. 2024. [JetMoE: Reaching Llama2 Performance with 0.1M Dollars.](#)
- [157] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2020. [Megatron-LM: Training Multi-Billion Parameter Language Models Using Model Parallelism.](#)
- [158] Shivalika Singh, Freddie Vargus, Daniel Dsouza, Börje F. Karlsson, Abinaya Mahendiran, Wei-Yin Ko, Herumb Shandilya, Jay Patel, Deividas Mataciunas, Laura OMahony, Mike Zhang, Ramith Hettiarachchi, Joseph Wilson, Marina Machado, Luisa Souza Moura, Dominik Krzemiński, Hakimeh Fadaei, Irem Ergün, Ifeoma Okoh, Aisha Alaagib, Oshan Mudannayake, Zaid Alyafeai, Vu Minh Chien, Sebastian Ruder, Surya Guthikonda, Emad A. Alghamdi, Sebastian Gehrmann, Niklas Muennighoff, Max Bartolo, Julia Kreutzer, Ahmet Üstün, Marzieh Fadaee, and Sara Hooker. 2024. [Aya Dataset: An Open-Access Collection for Multilingual Instruction Tuning.](#)

- [159] Snowflake. 2024. [Snowflake Arctic Cookbook Series: Exploring Mixture of Experts \(MoE\)](#).
- [160] Snowflake. 2024. [Snowflake Arctic: The Best LLM for Enterprise AI — Efficiently Intelligent, Truly Open](#).
- [161] Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Authur, Ben Bogin, Khyathi Chandu, Jennifer Dumas, Yanai Elazar, Valentin Hofmann, Ananya Harsh Jha, Sachin Kumar, Li Lucy, Xinxu Lyu, Nathan Lambert, Ian Magnusson, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Abhilasha Ravichander, Kyle Richardson, Zejiang Shen, Emma Strubell, Nishant Subramani, Oyvind Tafjord, Pete Walsh, Luke Zettlemoyer, Noah A. Smith, Hannaneh Hajishirzi, Iz Beltagy, Dirk Groeneveld, Jesse Dodge, and Kyle Lo. 2024. [Dolma: an Open Corpus of Three Trillion Tokens for Language Model Pretraining Research](#).
- [162] Luca Soldaini and Kyle Lo. 2023. [peS2o \(Pretraining Efficiently on S2ORC\) Dataset](#).
- [163] Guijin Son, Hanwool Lee, Sungdong Kim, Seungone Kim, Niklas Muennighoff, Taekyoon Choi, Cheonbok Park, Kang Min Yoo, and Stella Biderman. 2024. [KMMLU: Measuring Massive Multitask Language Understanding in Korean](#).
- [164] Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. 2023. [RoFormer: Enhanced Transformer with Rotary Position Embedding](#).
- [165] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2020. [VL-BERT: Pre-training of Generic Visual-Linguistic Representations](#).
- [166] 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. 2024. [Branch-Train-MiX: Mixing Expert LLMs into a Mixture-of-Experts LLM](#).
- [167] Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. 2022. [Challenging BIG-Bench Tasks and Whether Chain-of-Thought Can Solve Them](#).
- [168] Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge](#).
- [169] Shawn Tan, Yikang Shen, Zhenfang Chen, Aaron Courville, and Chuang Gan. 2023. [Sparse Universal Transformer](#).
- [170] Chaofan Tao, Qian Liu, Longxu Dou, Niklas Muennighoff, Zhongwei Wan, Ping Luo, Min Lin, and Ngai Wong. 2024. [Scaling Laws with Vocabulary: Larger Models Deserve Larger Vocabularies](#).
- [171] Chameleon Team. 2024. [Chameleon: Mixed-Modal Early-Fusion Foundation Models](#).
- [172] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, et al. 2023. [Gemini: A Family of Highly Capable Multimodal Models](#).
- [173] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, Soroosh Mariooryad, Yifan Ding, Xinyang Geng, Fred Alcober, Roy Frostig, Mark Omernick, Lexi Walker, Cosmin Paduraru, Christina Sorokin, Andrea Tacchetti, Colin Gaffney, Samira Daruki, Olcan Sercinoglu, Zach Gleicher, Juliette Love, Paul Voigtlaender, Rohan Jain, et al. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#).
- [174] Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Pier Giuseppe Sessa, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon, Machel Reid, Maciej Mikula, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain, Olivier Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Rahma Chaabouni, Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruiho Liu, Ryan Mullins, Samuel L Smith, Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya, Siamak Shakeri, Soham De, Ted Klimentko, Tom Hennigan, Vlad Feinberg, Wojciech Stokowiec, Yu hui

- Chen, Zafarali Ahmed, Zhitao Gong, Tris Warkentin, Ludovic Peran, Minh Giang, Clément Farabet, Oriol Vinyals, Jeff Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin Ghahramani, Douglas Eck, Joelle Barral, Fernando Pereira, Eli Collins, Armand Joulin, Noah Fiedel, Evan Senter, Alek Andreev, and Kathleen Kenealy. 2024. [Gemma: Open Models Based on Gemini Research and Technology](#).
- [175] Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, et al. 2024. [Gemma 2: Improving Open Language Models at a Practical Size](#).
- [176] Jamba Team, Barak Lenz, Alan Arazi, Amir Bergman, Avshalom Manevich, Barak Peleg, Ben Aviram, Chen Almagor, Clara Fridman, Dan Padnos, Daniel Gissin, Daniel Jannai, Dor Muhlgay, Dor Zimberg, Edden M Gerber, Elad Dolev, Eran Krakovsky, Erez Safahi, Erez Schwartz, Gal Cohen, Gal Shachaf, Haim Rozenblum, Hofit Bata, Ido Blass, Inbal Magar, Itay Dalmedigos, Jhonathan Osin, Julie Fadlon, Maria Rozman, Matan Danos, Michael Gokhman, Mor Zusman, Naama Gidron, Nir Ratner, Noam Gat, Noam Rozen, Oded Fried, Ohad Leshno, Omer Antverg, Omri Abend, Opher Lieber, Or Dagan, Orit Cohavi, Raz Alon, Ro'i Belson, Roi Cohen, Rom Gilad, Roman Glozman, Shahar Lev, Shaked Meirom, Tal Delbari, Tal Ness, Tomer Asida, Tom Ben Gal, Tom Braude, Uriya Pumerantz, Yehoshua Cohen, Yonatan Belinkov, Yuval Globerson, Yuval Peleg Levy, and Yoav Shoham. 2024. [Jamba-1.5: Hybrid Transformer-Mamba Models at Scale](#).
- [177] MosaicML NLP Team. 2023. [Introducing MPT-7B: A New Standard for Open-Source, Commercially Usable LLMs](#).
- [178] Qwen Team. 2024. [Qwen1.5-MoE: Matching 7B Model Performance with 1/3 Activated Parameters](#).
- [179] Reka Team, Aitor Ormazabal, Che Zheng, Cyprien de Masson d'Autume, Dani Yogatama, Deyu Fu, Donovan Ong, Eric Chen, Eugenie Lamprecht, Hai Pham, Isaac Ong, Kaloyan Aleksiev, Lei Li, Matthew Henderson, Max Bain, Mikel Artetxe, Nishant Relan, Piotr Padlewski, Qi Liu, Ren Chen, Samuel Phua, Yazheng Yang, Yi Tay, Yuqi Wang, Zhongkai Zhu, and Zhihui Xie. 2024. [Reka Core, Flash, and Edge: A Series of Powerful Multimodal Language Models](#).
- [180] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [LLaMA: Open and Efficient Foundation Language Models](#).
- [181] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open Foundation and Fine-Tuned Chat Models](#).
- [182] Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra, Clémentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M. Rush, and Thomas Wolf. 2023. [Zephyr: Direct Distillation of LM Alignment](#).
- [183] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2023. [Attention Is All You Need](#).
- [184] Ben Wang and Aran Komatsuzaki. 2021. [GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model](#).
- [185] Xingyao Wang, Boxuan Li, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, Jaskirat Singh, Hoang H. Tran, Fuqiang Li, Ren Ma, Mingzhang Zheng, Bill Qian, Yanjun Shao, Niklas Muennighoff, Yizhe Zhang, Binyuan Hui, Junyang Lin, Robert Brennan, Hao Peng, Heng Ji, and Graham Neubig. 2024. [OpenDevin: An Open Platform for AI Software Developers as Generalist Agents](#).
- [186] Yizhong Wang, Hamish Ivison, Pradeep Dasigi, Jack Hessel, Tushar Khot, Khyathi Raghavi Chandu, David Wadden, Kelsey MacMillan, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. 2023. [How Far Can Camels Go? Exploring the State of Instruction Tuning on Open Resources](#).

- [187] Zhilin Wang, Yi Dong, Olivier Delalleau, Jiaqi Zeng, Gerald Shen, Daniel Egert, Jimmy J. Zhang, Makesh Narsimhan Sreedhar, and Oleksii Kuchaiev. 2024. [HelpSteer2: Open-source dataset for training top-performing reward models](#).
- [188] Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. [Finetuned Language Models Are Zero-Shot Learners](#).
- [189] 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. 2024. [Skywork-MoE: A Deep Dive into Training Techniques for Mixture-of-Experts Language Models](#).
- [190] Johannes Welbl, Nelson F. Liu, and Matt Gardner. 2017. [Crowdsourcing Multiple Choice Science Questions](#).
- [191] BigScience Workshop, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, et al. 2023. [BLOOM: A 176B-Parameter Open-Access Multilingual Language Model](#).
- [192] Jialin Wu, Xia Hu, Yaqing Wang, Bo Pang, and Radu Soricut. 2024. [Omni-SMoLA: Boosting Generalist Multimodal Models with Soft Mixture of Low-rank Experts](#).
- [193] Shaohua Wu, Jiangang Luo, Xi Chen, Lingjun Li, Xudong Zhao, Tong Yu, Chao Wang, Yue Wang, Fei Wang, Weixu Qiao, Houbo He, Zeru Zhang, Zeyu Sun, Junxiong Mao, and Chong Shen. 2024. [Yuan 2.0-M32: Mixture of Experts with Attention Router](#).
- [194] xAI. 2024. [Open Release of Grok-1](#).
- [195] Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. 2023. [C-Pack: Packaged Resources To Advance General Chinese Embedding](#).
- [196] Cheng Xu, Shuhao Guan, Derek Greene, and M-Tahar Kechadi. 2024. [Benchmark Data Contamination of Large Language Models: A Survey](#).
- [197] Fuzhao Xue, Zian Zheng, Yao Fu, Jinjie Ni, Zangwei Zheng, Wangchunshu Zhou, and Yang You. 2024. [OpenMoE: An Early Effort on Open Mixture-of-Experts Language Models](#).
- [198] Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, Fan Yang, Fei Deng, Feng Wang, Feng Liu, Guangwei Ai, Guosheng Dong, Haizhou Zhao, Hang Xu, Haoze Sun, Hongda Zhang, Hui Liu, Jiaming Ji, Jian Xie, JunTao Dai, Kun Fang, Lei Su, Liang Song, Lifeng Liu, Liyun Ru, Luyao Ma, Mang Wang, Mickel Liu, MingAn Lin, Nuolan Nie, Peidong Guo, Ruiyang Sun, Tao Zhang, Tianpeng Li, Tianyu Li, Wei Cheng, Weipeng Chen, Xiangrong Zeng, Xiaochuan Wang, Xiaoxi Chen, Xin Men, Xin Yu, Xuehai Pan, Yanjun Shen, Yiding Wang, Yiyu Li, Youxin Jiang, Yuchen Gao, Yupeng Zhang, Zenan Zhou, and Zhiying Wu. 2023. [Baichuan 2: Open Large-scale Language Models](#).
- [199] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. 2024. [Qwen2 Technical Report](#).
- [200] John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. 2024. [SWE-agent: Agent-Computer Interfaces Enable Automated Software Engineering](#).
- [201] Zheng-Xin Yong, Hailey Schoelkopf, Niklas Muennighoff, Alham Fikri Aji, David Ifeoluwa Adelani, Khalid Almubarak, M Saiful Bari, Lintang Sutawika, Jungo Kasai, Ahmed Baruwa, Genta Indra Winata, Stella Biderman, Edward Raff, Dragomir Radev, and Vassilina Nikoulina. 2023. [BLOOM+1: Adding Language Support to BLOOM for Zero-Shot Prompting](#).
- [202] Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T. Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. 2024. [MetaMath: Bootstrap Your Own Mathematical Questions for Large Language Models](#).

- [203] Longfei Yun, Yonghao Zhuang, Yao Fu, Eric P Xing, and Hao Zhang. 2024. [Toward Inference-optimal Mixture-of-Expert Large Language Models](#).
- [204] Ted Zadouri, Ahmet Üstün, Arash Ahmadian, Beyza Ermiş, Acyr Locatelli, and Sara Hooker. 2023. [Pushing Mixture of Experts to the Limit: Extremely Parameter Efficient MoE for Instruction Tuning](#).
- [205] Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a Machine Really Finish Your Sentence?](#)
- [206] Biao Zhang and Rico Sennrich. 2019. [Root Mean Square Layer Normalization](#).
- [207] Ge Zhang, Scott Qu, Jiaheng Liu, Chenchen Zhang, Chenghua Lin, Chou Leuang Yu, Danny Pan, Esther Cheng, Jie Liu, Qunshu Lin, Raven Yuan, Tuney Zheng, Wei Pang, Xinrun Du, Yiming Liang, Yinghao Ma, Yizhi Li, Ziyang Ma, Bill Lin, Emmanouil Benetos, Huan Yang, Juntao Zhou, Kaijing Ma, Minghao Liu, Morry Niu, Noah Wang, Quehry Que, Ruibo Liu, Sine Liu, Shawn Guo, Soren Gao, Wangchunshu Zhou, Xinyue Zhang, Yizhi Zhou, Yubo Wang, Yuelin Bai, Yuhang Zhang, Yuxiang Zhang, Zenith Wang, Zhenzhu Yang, Zijian Zhao, Jiajun Zhang, Wanli Ouyang, Wenhao Huang, and Wenhao Chen. 2024. [MAP-Neo: Highly Capable and Transparent Bilingual Large Language Model Series](#).
- [208] Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. 2024. [TinyLlama: An Open-Source Small Language Model](#).
- [209] Qizhen Zhang, Nikolas Gritsch, Dwaraknath Gnaneshwar, Simon Guo, David Cairuz, Bharat Venkitesh, Jakob Foerster, Phil Blunsom, Sebastian Ruder, Ahmet Ustun, and Acyr Locatelli. 2024. [BAM! Just Like That: Simple and Efficient Parameter Upcycling for Mixture of Experts](#).
- [210] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. [OPT: Open Pre-trained Transformer Language Models](#).
- [211] Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, Less Wright, Hamid Shojanazeri, Myle Ott, Sam Shleifer, Alban Desmaison, Can Balioglu, Pritam Damania, Bernard Nguyen, Geeta Chauhan, Yuchen Hao, Ajit Mathews, and Shen Li. 2023. [PyTorch FSDP: Experiences on Scaling Fully Sharded Data Parallel](#).
- [212] Tianyu Zheng, Ge Zhang, Tianhao Shen, Xueling Liu, Bill Yuchen Lin, Jie Fu, Wenhao Chen, and Xiang Yue. 2024. [Opencodeinterpreter: Integrating code generation with execution and refinement](#). *arXiv preprint arXiv:2402.14658*.
- [213] Zexuan Zhong, Mengzhou Xia, Danqi Chen, and Mike Lewis. 2024. [Lory: Fully Differentiable Mixture-of-Experts for Autoregressive Language Model Pre-training](#).
- [214] Chunting Zhou, Pengfei Liu, Puxin Xu, Srini Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. 2023. [LIMA: Less Is More for Alignment](#).
- [215] Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. [Instruction-Following Evaluation for Large Language Models](#).
- [216] Yanqi Zhou, Nan Du, Yanping Huang, Daiyi Peng, Chang Lan, Da Huang, Siamak Shakeri, David So, Andrew Dai, Yifeng Lu, Zhifeng Chen, Quoc Le, Claire Cui, James Laudon, and Jeff Dean. 2024. [Brainformers: Trading Simplicity for Efficiency](#).
- [217] Yanqi Zhou, Tao Lei, Hanxiao Liu, Nan Du, Yanping Huang, Vincent Zhao, Andrew Dai, Zhifeng Chen, Quoc Le, and James Laudon. 2022. [Mixture-of-Experts with Expert Choice Routing](#).
- [218] Terry Yue Zhuo, Armel Zebaze, Nitchakarn Suppattarachai, Leandro von Werra, Harm de Vries, Qian Liu, and Niklas Muennighoff. 2024. [Astraios: Parameter-Efficient Instruction Tuning Code Large Language Models](#).
- [219] Barret Zoph, Irwan Bello, Sameer Kumar, Nan Du, Yanping Huang, Jeff Dean, Noam Shazeer, and William Fedus. 2022. [ST-MoE: Designing Stable and Transferable Sparse Expert Models](#).
- [220] Simiao Zuo, Xiaodong Liu, Jian Jiao, Young Jin Kim, Hany Hassan, Ruofei Zhang, Tuo Zhao, and Jianfeng Gao. 2022. [Taming Sparsely Activated Transformer with Stochastic Experts](#).
- [221] Ahmet Üstün, Viraat Aryabumi, Zheng-Xin Yong, Wei-Yin Ko, Daniel D’souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, Freddie Vargus, Phil Blunsom, Shayne Longpre, Niklas Muennighoff, Marzieh Fadaee, Julia Kreutzer, and Sara Hooker. 2024. [Aya Model: An Instruction Finetuned Open-Access Multilingual Language Model](#).

1	Introduction	2
2	Pretraining and adaptation	2
3	Results	4
4	MoE analysis	5
4.1	Router saturation	5
4.2	Expert co-activation	6
4.3	Domain specialization	6
4.4	Vocabulary specialization	7
5	Conclusion	8
A	Related work	26
B	Experimenting with alternative design choices	26
B.1	MoE-specific pretraining settings	26
B.1.1	Mixture-of-Experts vs. dense	26
B.1.2	Expert granularity	27
B.1.3	Shared experts	28
B.1.4	Expert choice vs. token choice	29
B.1.5	Sparse upcycling	29
B.1.6	Load balancing loss	30
B.1.7	Router z-loss	31
B.2	General pretraining settings	32
B.2.1	Dataset experiments	32
B.2.2	Initialization	32
B.2.3	RMSNorm	33
B.2.4	Decaying embedding parameters	34
B.2.5	QK-Norm	34
B.2.6	AdamW epsilon	34
B.3	Adaptation settings	35
C	Training configuration	36
D	Evaluation setup	40
E	Openness of models	41
F	Additional evaluation	43
G	Additional experiments	47
H	Analysis	49
H.1	Details of analysis in §4	49
H.2	Additional Analysis	50
I	Artifacts	58
J	Limitations and future work	58

A Related work

Advances in MoEs Current LMs still largely follow the transformer architecture [183] with only few architectural changes that have been widely adopted, such as decoder-only training [135], SwiGLU activations [151, 41], RoPE [164], MQA/GQA [150, 3] and RMSNorm [206]. Model sparsity via Mixture-of-Experts is one modification still under active exploration with some early adoption but most LMs, including Llama 3 [50], still rely on a dense architecture. There has been a lot of progress in improving the sparsely-gated MoE layer since its introduction [152]: New routing techniques [88, 144, 220, 66, 77, 49, 213, 193, 123], fine-grained expert segmentation [39, 69], stability [219] and efficiency [87, 139, 48, 216, 90, 166, 127, 143] improvements. In this work, we perform many experiments to provide insights into training Mixture-of-Experts LMs. Subsequently, we train **OLMoE-1B-7B** for 5T tokens. No prior MoE has been overtrained [57] to this extent to our knowledge making **OLMoE-1B-7B** the best testbed to research performance saturation of MoEs vs. dense models. With **OLMoE** we hope to facilitate such and other research to help the field uncover whether MoEs should make it into all future LMs and with what precise configuration.

Open LMs A variety of model families have been proposed under varying degrees of openness commonly categorized based on whether model weights are available. **Closed-weight** models include GPT [24, 126], Gemini [172, 173], PaLM [30, 9], Reka [179], and **open-weight** ones include Llama [180, 181, 50], Mistral [78, 79], Gemma [174, 175], Falcon [8, 130], MPT [177], Qwen [13, 199], GLM [61], Yi [2], DeepSeek [42, 43, 39], Nemotron [128, 125, 187], InternLM [26], Baichuan [198], Phi [68, 93, 1], StableLM [16], OPT [210], Zamba [62]. However, besides model weights, training data and code are key to enabling scientific research of these models [104, 105] and distributing their benefits broadly [23]. There have been few releases also including data and code in addition to model weights which we refer to as “**fully open-source**”: BLOOM [191, 149, 122, 201], GPT-NeoX [21, 22, 184], StarCoder [91, 108, 5, 119, 218], Pythia [18], OLMo [65], LLM360 [102], Cerebras-GPT [46], DCLM [89], MAP-Neo [207], RWKV [131, 132], and SmoLM [6]. For Mixture-of-Experts only OpenMoE [197] aims to be fully open-source, however, its poor performance limits its usefulness. We release **OLMoE-1B-7B** as the first state-of-the-art Mixture-of-Experts LM that is fully open-source: model weights, data, code, and logs.

B Experimenting with alternative design choices

In this section, we present pretraining and adaptation experiments that have led to **OLMoE-1B-7B**. We group them into experiments on settings specific to Mixture-of-Experts (§B.1), experiments on settings applicable to both dense LMs and MoEs (§B.2), and adaptation experiments (§B.3). In pretraining experiments, we often use MMLU Var, a version of MMLU [70] with varying few-shots and a different format that provides signal earlier during training. We describe our full evaluation setup in §D and provide additional experiments in §G. Each experiment links to a Weights & Biases report with more validation and downstream results, and the full configurations of the runs. To isolate the impact of changes and minimize confounders, we vary only one hyperparameter for each experiment. Nevertheless, due to the large number of hyperparameters, some results may change under different configurations and we cannot guarantee the correctness of each of our hyperparameter choices. Models are not comparable across different experiments, as we vary the base model to incorporate successful findings.

B.1 MoE-specific pretraining settings

B.1.1 Mixture-of-Experts vs. dense

Prior work reports various speed-ups of MoEs over dense models: Artetxe et al. [10] report that MoEs require 2–4× less compute to match dense models, MoMa [99] exhibits 2.6× FLOP savings for language tasks, Arctic [159] yields 4× FLOP savings but for very different dense and MoE configurations, and Switch Transformers [56] train 2–7× faster with MoEs but for encoder-decoder models while the other works study decoder-only LMs [135].

In Figure B1, we compare MoEs and dense models in a controlled setup. We find that our MoE reaches the performance of the dense model with ~3× fewer tokens equivalent to ~3× less compute measured in FLOPs. However, due to the additional memory overhead of training the MoE with

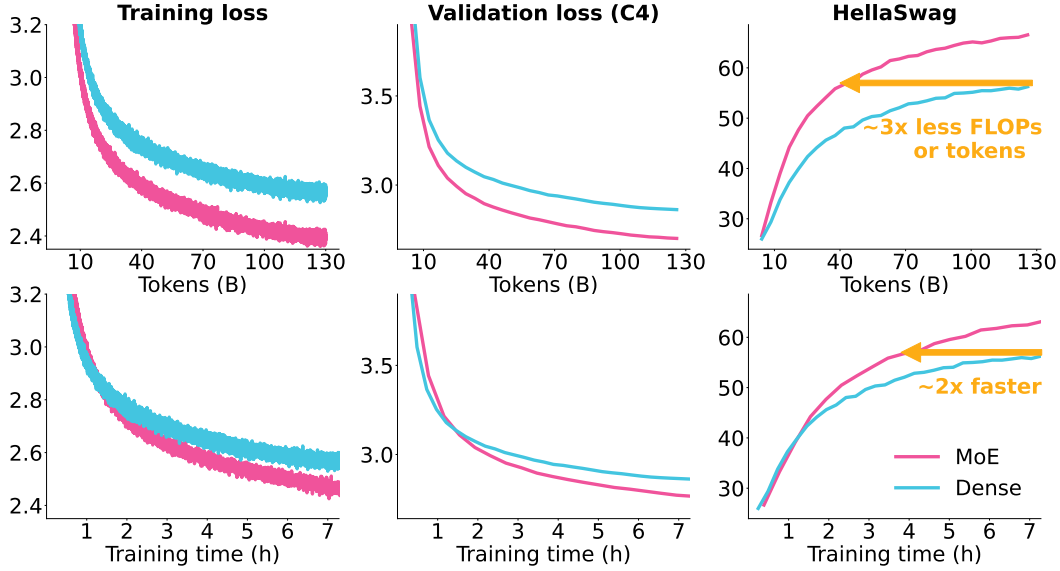


Figure B1: **MoE vs. Dense.** We train a 1.3B parameter dense model and a 1.3B active, 6.9B total parameter MoE model, each on 128 H100 GPUs. Apart from MoE-related changes, we train both with the same configuration for 130B tokens. The MoE contains 64 experts out of which 8 are activated with an FFN dimension of 1,024, while the dense model has an FFN dimension of 8,192. Thus both have the same number of active parameters. **Top:** The MoE reaches the final dense performance with $\sim 3\times$ fewer tokens (or FLOPs, as both have the same active parameters ignoring the trivial router parameters). **Bottom:** Due to some memory overhead, this equates to $\sim 2\times$ faster training. More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-MoE-vs-Dense--Vml1dzo40TMOMjkx>

its 7B total parameters, it processes fewer tokens per second than the dense model (23,600 tokens per second per GPU for the MoE vs. 37,500 for dense). Thus, in terms of training time, it reaches the performance of the dense model only $\sim 2\times$ faster. There are likely optimizations possible that would bring the speed-up closer to the $3\times$ token speed-up, which we leave to future work. Based on these results, we select an MoE configuration with 6.9B total and 1.3B active parameters matching OLMo-7B in total and OLMo-1B in active parameter count, respectively.

B.1.2 Expert granularity

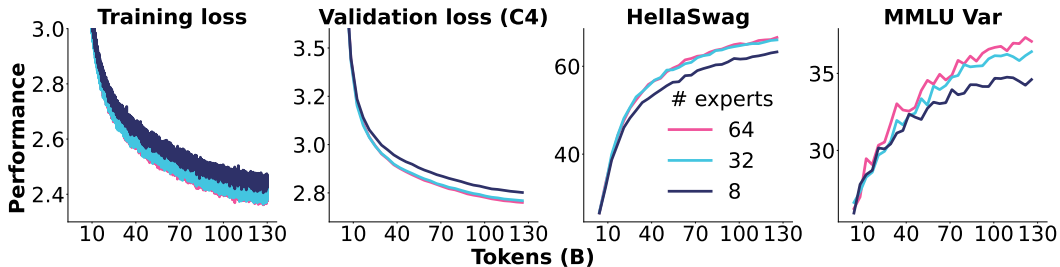


Figure B2: **Expert granularity.** We vary the number of experts in tandem with the FFN dimension to ensure that active and total parameters and thus compute cost remain the same. For example, for 64 experts, the FFN dimension is 1,024 and 8 experts are activated, while for 32 experts it is 2,048 with 4 activated experts. More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-Granularity--Vml1dzo40Tlx0TE4>

Dai et al. [39] propose to use small fine-grained experts to allow more combinations of experts and thus make the model more flexible. For example, the Mixtral model [79] uses the common configuration of 8 experts per layer, 2 of which are activated. This allows for $\binom{8}{2} = 28$ combinations per layer. By halving the size of each expert and therefore doubling the number of experts to maintain the same compute and parameter budget, we can increase the possible combinations to $\binom{16}{4} = 1,820$. Krajewski et al. [86] investigate compute-optimal granularity configurations finding that higher compute budgets warrant more granular experts.

In Figure B2, we observe that more granular experts improve training loss, validation loss, and downstream performance. The 8-expert configuration uses 1 active expert, which yields $\binom{8}{1} = 8$ combinations. By quartering the size of each expert but increasing the number to 32 with 4 active ones ($\binom{32}{4} = 35,960$ combinations), we observe an improvement of around 10% on HellaSwag and MMLU at around 130 billion tokens. However, we find that there are diminishing returns to granularity. The additional increase to 64 experts with 8 active ones ($\binom{64}{8} = 4,426,165,368$ combinations) improves downstream metrics by a smaller amount of 1–2%. For our **OLMOE-1B-7B** compute budget³ of 3×10^{22} , Krajewski et al. [86] predict an optimal number of experts of 256 ($G = 32$ in their paper). However, their predictions are for compute-optimal models [72, 32], while we train for 5T tokens, which is orders of magnitude beyond what would be conventionally considered optimal for our model size. Thus, their predictions may not extend to our setup, and we stick with 64 experts for **OLMOE-1B-7B**, also due to the diminishing returns in Figure B2.

B.1.3 Shared experts

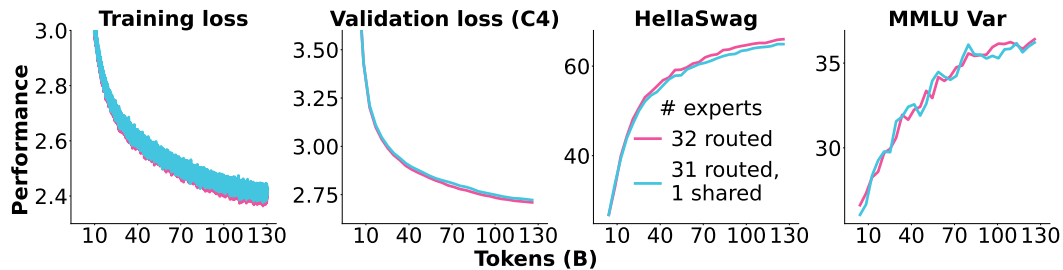


Figure B3: **Shared experts**. Both setups have the same number of active and total parameters and use the same number of FLOPs. 4 of the 32 routed experts are activated, while it is 3 for the 31 routed experts of the other model, as it has 1 always-active shared expert. More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-Expert-sharing--Vmlldzo40TIyMjQz>

Dai et al. [39] propose training with a shared/fixe expert that is always used in addition to the routed experts. The intuition is to encourage the shared expert to learn common information and allow the other routed experts to learn more specialized knowledge. This should reduce redundancy among experts and thus lead to a better model as it can store more total information.

In Figure B3, we benchmark having a single shared and a single routed expert versus two routed experts. While both settings lead to similar performance, sharing an expert performs slightly worse. Sharing an expert removes flexibility from the model and thus goes against the findings in §B.1.2 suggesting that allowing for more expert combinations improves performance. Specifically, the two models in Figure B3 have $\binom{32}{4} = 35,960$ and $\binom{31}{3} = 4,495$ possible combinations per layer. Thus, removing one of the routed experts and turning it into a shared one eliminates almost 90% of possible combinations. This likely acts as a counterforce to the potential benefits of isolating common knowledge in a shared expert. Based on these results, we do not use shared experts in **OLMOE-1B-7B**, but we do think that there is merit to the idea of experts that are activated more often or even always. However, rather than enforcing this behavior via a shared expert, we believe that it should be learned

³Approximated via $6 * N * D$ [80], where N are active parameters (1B) and D are training tokens (5T).

by the model. This is difficult with current setups due to the necessity of a load balancing loss (§B.1.6) penalizing the model if tokens are not distributed equally among experts. Potential future work can explore removing the load balancing loss to allow for more flexible usage of experts.

B.1.4 Expert choice vs. token choice




Figure B4: **Expert choice (EC) vs. token choice (TC)**. Both models have an 8-expert MoE in every 2nd layer. For TC, 2 experts are activated per token, while for EC the capacity factor is 2. Thus, both models use the same number of active parameters. More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-EC-vs-TC--Vml1dzo4MzkzMDM3>

The MoE router determines which experts process each input token (§2). There are two common types [101]: **expert choice (EC)** [217] and **token choice (TC)** [152]. For EC, each expert selects a fixed number of tokens from the incoming sequence. By design, this leads to each expert processing the same number of tokens. This is the main benefit of EC as it ensures perfect load balance, which improves training throughput and removes the need for a load balancing loss. The main downside of EC is that it is not easily usable for autoregressive generation where a single token is processed at each step rather than the entire sequence in one [141]. Another potential downside is that EC can lead to token dropping, where some tokens are not selected by any expert, which can hurt performance [58]. At the same time, it can lead to some tokens being processed by multiple experts, which could also be beneficial as it allows the model to allocate more compute to some tokens [217]. For TC, each token selects a fixed number of experts. This can lead to many tokens choosing the same expert, hurting training efficiency. Therefore it is common to use TC with a load balancing loss [152] to encourage equal distribution.

In Figure B4, we benchmark EC and TC. We find that TC outperforms EC for the same token budget for all tasks depicted as well as other tasks like PIQA, SciQ, etc. which we report at <https://wandb.ai/ai2-llm/olmoe/reports/Plot-EC-vs-TC--Vml1dzo4MzkzMDM3>. While Zhou et al. [217] find EC to be better, our configuration slightly differs in that we use dropless MoEs [58] with a load balancing loss. Thus, our TC variant is expected to perform better than the TC variant in Zhou et al. [217]. We confirm findings that EC runs around 20% faster at 29,400 tokens per second per device versus 24,400 for TC [217]. EC may be more beneficial in a multimodal setup [99] as dropping noisy image tokens is likely less harmful than text tokens. Thus, while we stick with TC for this release of **OLMoE**, we may revisit EC for future multimodal models.

B.1.5 Sparse upcycling

Komatsuzaki et al. [85] propose turning a dense model into a Mixture-of-Experts model via sparse upcycling: (1) The dense MLP is cloned for each desired expert to constitute MoE layers. (2) A newly initialized router is added in front of each MoE layer. (3) Pretraining continues with the new model so that the cloned MLPs can gradually specialize in different things and the router can be learned. They find that the upcycling approach maintains a performance advantage over a language model trained from scratch for up to 120% of the compute budget of the original dense checkpoint that the sparse model was upcycled from. For example, if sparsely upcycling a 1.3B parameter model at 2 trillion tokens then only at 2.4 trillion tokens should an MoE trained from scratch catch up with the upcycled model. That is, the sparsely upcycled model would have been trained for another 400 billion tokens, thereby saving the equivalent of up to 2T tokens of compute. Other works such as MiniCPM [74], Qwen2 [199] and reportedly Mixtral [25, 79] have adopted sparse upcycling but only share limited information about their configuration.

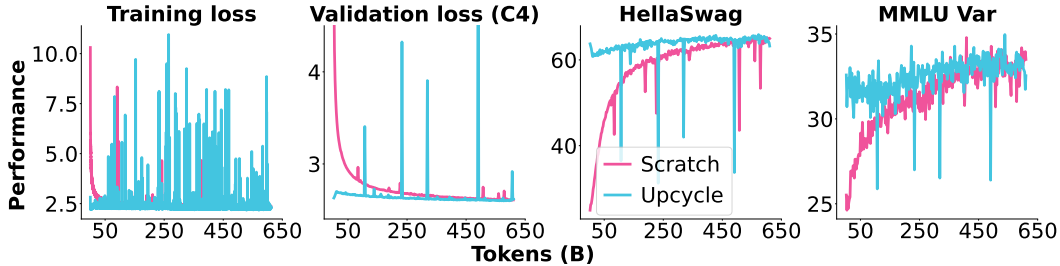


Figure B5: **Sparse upcycling.** We upcycle OLMo-1B (0724) at 2T tokens into an MoE with 8 total experts of which 2 are activated and train it for an additional 610 billion tokens. We compare it to a model trained from scratch for 610 billion tokens. Except for this difference, both models use the same config, which includes some suboptimal settings that contribute to the instability, such as no QK-Norm (§B.2.5) and no truncated normal init (§B.2.2). More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-Scratch-vs-Upcycle--Vml1dz04NDIyOTc4>

In Figure B5, we compare sparse upcycling OLMo-1B (0724) [65] with training an MoE from scratch. We find that after 500B tokens, an otherwise equivalent MoE trained from scratch already catches up with the upcycled model both on the metrics in Figure B5 and our additional metrics at <https://wandb.ai/ai2-llm/olmoe/reports/Plot-Scratch-vs-Upcycle--Vml1dz04NDIyOTc4>. At around 600B tokens, the MoE from scratch starts outperforming the upcycled MoE. Thus, it only requires 25% of the compute budget of the original dense model to catch up as opposed to the 120% reported in Komatsuzaki et al. [85]. However, they use expert choice routing and study encoder-decoder models [137]. Meanwhile, we use token choice routing (§B.1.4) and decoder-only models (§2). Further, we upcycle a model that has already been significantly overtrained [57], i.e., a 1B model trained for 2T tokens. Its parameters are likely already in a very optimal range for a dense model, which may limit the amount of additional exploration possible after upcycling. This motivates us to experiment with adding noise to the upcycled weights outlined in §G, but we do not find it to lead to better performance. A large disadvantage of upcycling is that the upcycled MoE is constrained by some hyperparameters of the dense model. Specifically, OLMo-1B (0724) was trained without QK-Norm and normal initialization, both of which hurt stability in our experiments (§B.2.5, §B.2.2). While it may be possible to simply add new QK-Norms and train them from scratch similar to the new router layer trained from scratch, it is impossible to change the initialization of the original dense model when upcycling it. Thus, as we want to change these hyperparameters and also train **OLMoE-1B-7B** for around 250% of the compute budget of the dense model (5T vs. 2T tokens), we do not use upcycling.

B.1.6 Load balancing loss




Figure B6: **Impact of applying a load balancing loss (LBL).** The training loss plot excludes the load balancing loss for both models. More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-LBL-vs-No-LBL--Vml1dz04OTkyNDg4>

Shazeer et al. [152] propose the load balancing loss to penalize the model if it is unbalanced, i.e., if it routes all tokens to only a few experts. This is based on the observation that without such penalty, models tend to update only a select few experts in each layer [52, 17]. To compute the load




Figure B7: **Expert assignment during training when using or not using a load balancing loss for the first MoE layer.** More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmo/reports/Plot-LBL-vs-No-LBL--Vml1dzo40TkyNDg4>

balancing loss (\mathcal{L}_{LB}) we multiply the fraction of tokens f_i routed to one expert E_i with the total routing probability P_i allocated to E_i for one batch and sum it across the number of experts N_E :

$$\mathcal{L}_{LB} = N_E \cdot \sum_{i=1}^{N_E} f_i \cdot P_i \quad (3)$$

The loss is further scaled by N_E and a loss weight α (see Equation 2), which is an optional weight to determine the magnitude of the loss commonly set to 0.01 [219, 197]. We do not experiment with changing the weight of 0.01.

In Figure B6 we investigate the performance impact of using the auxiliary load balancing loss. We find that across training loss and validation losses, using the load balancing loss leads to better performance even after only a few billion tokens. We still measure the load balancing loss even when it is not used (“No LBL”) and find that while it spikes initially, it slowly decreases over the next few billion tokens. This behavior is also visible in Figure B7 (left), where initially all tokens in the first layer are assigned to the 6th expert (pink). Eventually, the model also starts assigning some tokens to the 1st expert (yellow). However, all other experts remain largely flat and are thus “dead weights” that take up GPU memory but are not used. Given these results, we use the auxiliary load balancing loss with a weight of 0.01 following prior work [152, 156]. However, getting rid of the load balancing loss is an important direction for future research as it constrains the flexibility of the model by forcing it to use all experts approximately equally. This could prevent the experts from specializing in certain data domains and may be a reason prior work has failed to find strong evidence of expert specialization [79, 219].

B.1.7 Router z-loss

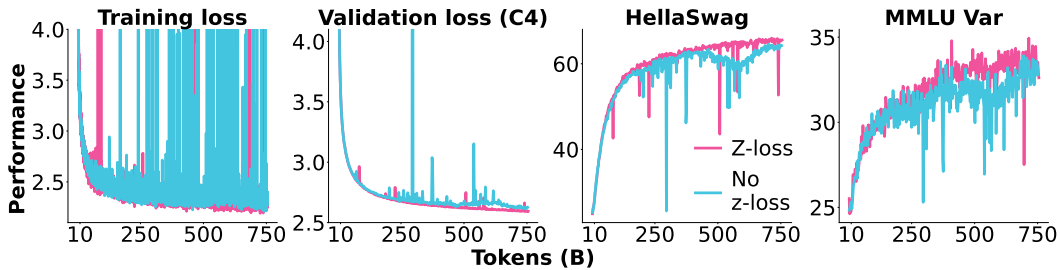


Figure B8: **Router z-loss.** We compare adding router z-loss with a loss weight of 0.001 versus no additional z-loss. More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmo/reports/Plot-Zloss-vs-none--Vml1dzo4NDM4NjUz>

Zoph et al. [219] propose the router z-loss to improve both the stability and quality of MoE models. This auxiliary loss penalizes large logits coming into the gating network. Such large logits can lead to numeric overflows in the large matrix multiplications happening in the MoE layer. It is computed

by exponentiating the logits x_j right before the router layer summed across the number of experts N_E and averaged across the batch B , thereby making larger logits lead to a larger loss:

$$\mathcal{L}_{RZ}(x) = \frac{1}{B} \cdot \sum_{i=1}^B \left(\log \sum_{j=1}^{N_E} \exp(x_j^{(i)}) \right)^2 \quad (4)$$

The loss is further multiplied with an optional loss weight, β (see Equation 2), to determine the magnitude of the loss commonly set to 0.001 [219, 156]. We do not experiment with changing the weight of 0.001.

In Figure B8, we confirm that across training loss, validation loss, and downstream performance adding the router z-loss improves stability (less spikes) and quality (lower loss and higher downstream performance). Thus, despite it reducing throughput by $\sim 2\%$ we use the router z-loss for **OLMoE-1B-7B** with a weight of 0.001 as in Zoph et al. [219].

B.2 General pretraining settings

B.2.1 Dataset experiments




Figure B9: **OLMoE-MIX vs. Dolma 1.7.** We compare our data mix described in §2 with Dolma 1.7 used to train prior OLMo models. Lower training loss does not mean that one dataset is better, but rather suggests which dataset is easier for the model to learn. More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-Dolma-1-7-vs-Dolma-OLMoE--Vmlldzo40TIxNTg5>

Li et al. [89] release the DCLM-Baseline dataset and establish that it leads to better language models than Dolma 1.7 and other datasets as measured on common benchmarks like MMLU [70]. This motivates us to mix their DCLM dataset with some components from Dolma 1.7 that we deem to be high-quality; see §2. In Figure B9, we compare our mix, **OLMoE-MIX**, with Dolma 1.7 in a controlled setup. We find that **OLMoE-MIX** leads to clear gains on all three downstream metrics, especially MMLU. DCLM-Baseline has been created through a series of dataset ablations targeting MMLU and other downstream metrics, which explains these results. We also compare adding Reddit and FLAN to our mix as detailed in §G, but do not find consistent performance gains. We do not have a strong intuition for why adding these datasets does not help and a more automatic approach to dataset mixing may be desirable for future iterations [100, 4].

We pretrain using our mix of DCLM-Baseline and Dolma 1.7 dubbed **OLMoE-MIX**.

B.2.2 Initialization

Few prior works on Mixture-of-Experts share their initialization strategy. Even the most open MoEs prior to this work, JetMoE [156] and OpenMoE [197], do not mention their initialization scheme. For DeepSeekMoE [39] and DeepSeekV2 [43], the authors share that they use a normal initialization with a standard deviation (std) of 0.006. For dense language models, a normal initialization with an std of 0.02 has been commonly used as popularized by Shoeybi et al. [157].

In Figure B10, we find a truncated normal initialization leads to more stable training and better performance than a regular normal initialization. The difference between the two initializations only becomes clear at around 450 billion tokens, where the model with the normal initialization starts to diverge. This is despite both models using the same configuration except for

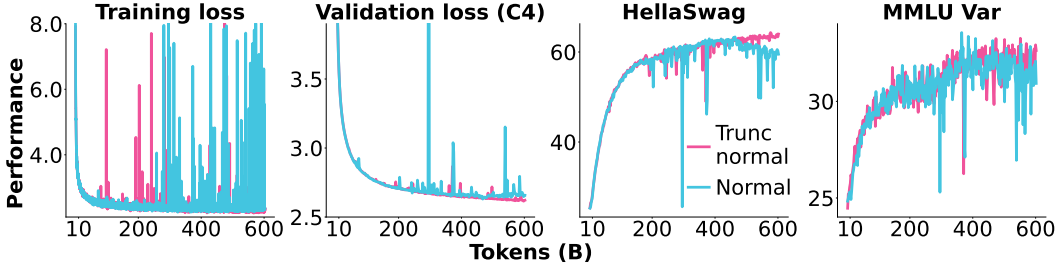


Figure B10: **Initialization.** We compare a normal initialization with a standard deviation (std) of 0.02 with a truncated normal initialization with a maximum (minimum) cut-off of 0.06 (-0.06) corresponding to three stds (3×0.02). More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-Init--Vmlldzo4NDIzMzM5>

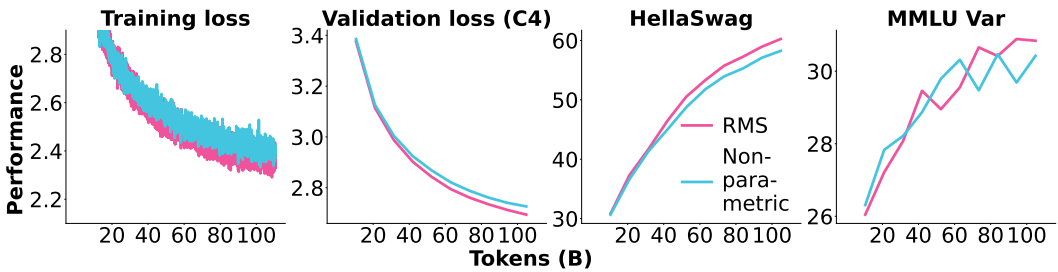


Figure B11: **Non-parametric layer normalization vs. RMSNorm.** More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-LN--Vmlldzo4NDQyMTAz>

the difference in weight initialization. Having to train for hundreds of billions of tokens until an experiment provides a clear signal is one of the key challenges of pretraining ablations.

We use the truncated normal initialization for **OLMoE-1B-7B**.

B.2.3 RMSNorm

OLMo [65] uses non-parametric layer normalization [12], mainly as it is significantly faster than the commonly used RMSNorm [206, 112]. This is an unusual choice as most LMs use RMSNorm, such as the Llama [180, 181, 50], Gemma [174, 175], and Qwen [13, 199] model families.

In Figure B11, we observe that replacing the non-parametric layer normalization in OLMo with a parametric RMSNorm leads to better performance. This is likely because the non-parametric layer normalization leads to a large number of spikes in the gradients as seen in Figure B13. We clip gradients at 1.0, which prevents these spikes from leading to very large and potentially disruptive parameter updates. However, the clipped gradients may still harm the performance of the model as they are no longer the true gradients. Thus, despite RMSNorm lowering our training throughput by 15%, we train our final model with RMSNorm. We include the RMSNorm parameters in weight decay as we find that it performs slightly better (Figure B12) even though it is common practice to exclude them.⁴




Figure B13: **Total norm of the gradients when training with RMS or non-parametric normalization.** We increase the logging interval of the RMS run at 75B tokens, hence its change in thickness.




Figure B12: **Decaying the RMSNorm parameters.** More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmo/reports/Plot-Decay-LN--Vml1dzo4NDQ1NDYy>




Figure B14: **Decaying the embedding parameters.** More results, logs, and configurations: <https://api.wandb.ai/links/ai2-llm/3h22onp5>

B.2.4 Decaying embedding parameters

Similar to the RMSNorm parameters (§B.2.3), embedding parameters are commonly excluded from weight decay.⁵ In Figure B14 we find that whether or not they are decayed has only a minor impact on performance, with decaying being slightly better. Thus for simplicity, we weight decay all parameters in **OLMOE-1B-7B** including embedding and RMSNorm.

B.2.5 QK-Norm

Some works have reported stability improvements from adding layer normalization after the query and key projections (“QK-Norm”) [171, 112, 44]. QK-Norm can prevent the subsequent attention operation from leading to very large logits that may lead to numeric overflows and destabilize the network, especially when training in low precision. Like layer normalization at other places in the model, the QK-Norm could be non-parametric or use the parametric RMSNorm (§B.2.3).

In Figure B15, we compare using QK-Norm with no normalization after the query and key projections. We find that QK-Norm leads to some stability and performance improvements. We perform this experiment with non-parametric layer normalization as used in OLMo [65], while we used parametric RMS layer normalization [206] for **OLMOE-1B-7B** (§B.2.3). To ensure the benefit of QK-Norm is not an artifact of comparing with non-parametric layer normalization, we run another experiment with RMS layer normalization and still find QK-Norm to lead to slightly better training loss and to prevent a large grad norm spike.⁶ Thus, we use QK-Norm for **OLMOE-1B-7B** despite it reducing throughput by almost 10%.

B.2.6 AdamW epsilon

Groeneveld et al. [65] use an epsilon (“eps”) value of 1E-05 in the AdamW optimizer for training OLMo. A larger eps value leads to smaller steps of the optimizer but can be more stable [83].

⁴<https://github.com/karpathy/minGPT/pull/24#issuecomment-679316025>

⁵<https://github.com/karpathy/minGPT/pull/24#issuecomment-679316025>

⁶<https://wandb.ai/ai2-llm/olmo/reports/Plot-QKNorm-revisited--Vml1dzo4NTc2NTIz>

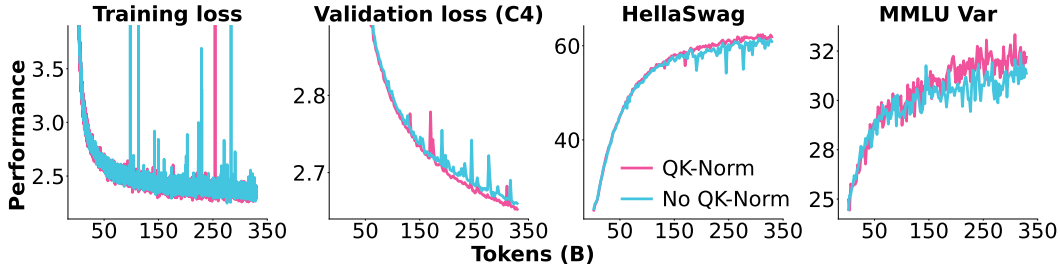


Figure B15: **Query-Key layer normalization (QK-Norm)**. Both models use non-parametric layer normalization. QK-Norm corresponds to additional layer normalization of the query and key projections. More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-QKNorm-vs-none--Vmlldzo4NDIzZmZlMzE2>




Figure B16: **AdamW epsilon**. More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-AdamW-eps--Vmlldzo4NDc5MDg0>

In Figure B16, we find that decreasing eps to the recommended default of 1E-08 [83] significantly improves performance while the run remains stable. Thus, we set eps to 1E-08 for our final run.

B.3 Adaptation settings

We experiment with small design choices for adaptation using our evaluation setup described in §D. (1) **Auxiliary losses:** Zoph et al. [219] find that using the auxiliary load balancing loss (§B.1.6) during regular finetuning leads to small performance gains. For instruction tuning, however, Shen et al. [154] do not find conclusive evidence in favor of using the load balancing or router z-loss with only small differences in performance, both in support of and against the auxiliary losses. In Table B2 we display experiments with the load balancing loss during adaptation and find that not using it leads to better performance (54.0 vs. 52.8 after instruction tuning (SFT) and 57.7 vs. 57.1 after preference tuning (DPO)). One potential problem of deactivating the load balancing loss is that it may harm balance among experts and turn some into dead weights as observed during pretraining in §B.1.6. However, when measuring the load balancing loss in Table B1 on our SFT data (§2), we find that the loss only increases by around 0.01% after SFT (8.250 vs. 8.249). This is likely because which experts certain tokens get routed to is determined early during pretraining, as we find later in the analysis section (§4.1). We also visualize the activation patterns of experts of the model after pretraining, and the models after SFT and DPO trained without load balancing in §H.2 (Figure H5) finding that the distribution remains around the same. Thus, as our models adapted without load balancing perform better and we find it not to impact routing substantially, we do not use load balancing during adaptation. (2) **Annealing checkpoint:**

Table B1: **Load balancing loss (Equation 3) over a subset of the respective corpora prior to scaling with the load balancing loss weight α** . While we use load balancing loss during pretraining, we do not use it during SFT.

Data (↓)	OLMOE-1B-7B	
	After pretraining	After SFT
Wikipedia	8.331	8.367
C4	8.073	8.076
SFT data	8.249	8.250

We also experiment with using the checkpoint pre-annealing (§2) for adaptation and find the checkpoint post-annealing leads to better performance (53.8 vs. 54.0 after SFT and 56.3 vs 57.7 after DPO), thus we use the post-annealing checkpoint. **(3) Preference algorithm:** Since the release of DPO (Direct Preference Optimization) [136], a variety of preference algorithms have been proposed [54, 73, 113]. We experiment with KTO [54] and find that it matches DPO in Table B2 for our setup (§C). While we release both models, we use DPO for our final **OLMoE-1B-7B-INSTRUCT** model, as it scores higher on AlpacaEval, which has a smaller chance of data contamination than our other benchmarks [196].

Table B2: **Adaptation experiments of OLMoE-1B-7B.** We compare using the pretrained checkpoint prior to annealing for adaptation, using the checkpoint after the additional 100B tokens of annealing, and using the checkpoint after the additional 100B tokens of annealing and with load balancing loss (§B.1.6) during adaptation. We apply DPO/KTO to the respective SFT model.

Task (→) Setup (→) Metric (→)	MMLU	GSM8k	BBH	Human-	Alpaca-	XSTest	IFEval	Avg
	0-shot	8-shot CoT	0-shot	Eval	Eval 1.0			
	EM	EM	EM	0-shot Pass@10	0-shot %win			
OLMoE-1B-7B w/o annealing	49.0	2.0	31.5	18.9	-	62.1	18.5	-
+SFT	50.2	43.0	35.6	55.5	68.9	83.8	39.7	53.8
+DPO	50.9	36.0	35.8	58.8	81.7	83.2	47.9	56.3
OLMoE-1B-7B +SFT	49.8	3.0	33.6	22.4	-	59.7	16.6	-
+SFT	51.4	40.5	38.0	51.6	69.2	84.1	43.3	54.0
+DPO	51.9	45.5	37.0	54.8	84.0	82.6	48.1	57.7
+KTO	51.2	45.5	34.1	57.1	81.6	86.6	47.5	57.7
+SFT (load balancing)	50.9	36.5	35.7	52.4	66.9	84.8	42.3	52.8
+DPO (load balancing)	51.1	42.5	39.3	55.6	82.9	82.1	46.0	57.1

C Training configuration

Table C1: **Composition of the pretraining data for OLMoE-1B-7B.** StarCoder, peS2o, and Wikipedia parts come from Dolma 1.7 [161]. Links to our data are in §I.

Source	Doc Type	GPT-NeoX tokens (billions)	Words (billions)	UTF-8 bytes (GB)	Documents (millions)
DCLM-Baseline [89]	web pages	3,860	3,380	16,700	2,950
StarCoder [91, 84]	code	101	63.9	325	78.7
peS2o [162, 161]	STEM papers	57.2	51.3	268	38.8
arXiv [36]	STEM papers	21.1	23.5	88.8	1.55
OpenWebMath [129]	math web pages	12.7	10.2	42.4	2.91
Algebraic Stack [11]	math proofs code	12.6	9.6	39.3	2.83
English Wikipedia & Wikibooks [161]	encyclopedic	3.69	3.16	16.2	6.17
Total		4,060	3,530	17,400	3,080

Pretraining We display the pretraining hyperparameter configuration of **OLMoE-1B-7B** in §C comparing with other relevant models. We follow Groeneveld et al. [65] using the AdamW optimizer [106] with ZeRO [140] via PyTorch FSDP [211] and mixed-precision training [115]. Our main model settings differing from Groeneveld et al. [65] are: **(1) MoE-related changes: OLMoE-1B-7B**

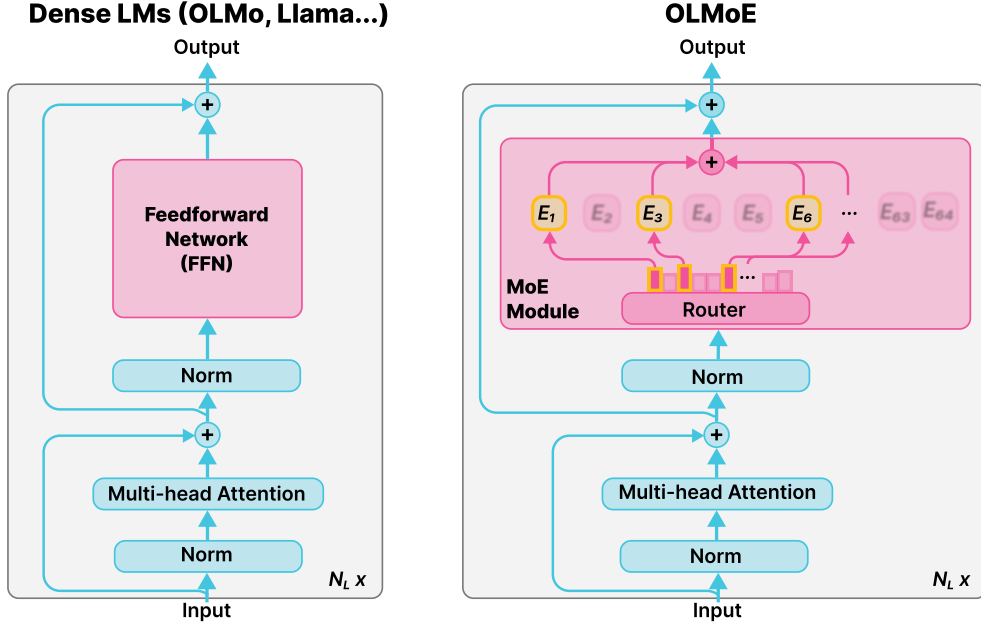


Figure C1: Comparison of the architecture of dense LMs and MoE models like OLMoE. The figure excludes some details, e.g., OLMoE-1B-7B also uses QK-Norm (§B.2.5).

Table C2: Adaptation training data for OLMoE-1B-7B. Links to our data are in §I.

Source	Domain	Samples
<i>Instruction Tuning</i>		
Tulu 2 SFT Mix [76]	Various	326,154
No Robots [138]	Various	9,500
CodeFeedback-Filtered-Instruction [212]	Coding	156,526
MetaMathQA [202]	Math	98,750
Advanced (non-chat) subset of Daring Anteater [187]	Various	17,082
<i>Preference Tuning (DPO [136])</i>		
UltraFeedback [38] binarized and filtered for TruthfulQA [98] contamination	Various	60,800

is a sparsely activated decoder-only transformer [183] using dropless Mixture-of-Experts [58]. Unlike most prior MoEs, we use a high granularity [39, 86] with 64 small experts with an FFN dimension of just 1,024 rather than a few large experts. We further use two auxiliary losses: router z-loss [219] and load balancing loss [152]. **(2) Stability improvements:** (a) We use a truncated normal initialization with a standard deviation of 0.02 and a minimum (maximum) cut-off of -0.06 (0.06) corresponding to three standard deviations. (b) We use QK normalization [171, 112, 44]. (c) We use RMSNorm [206] instead of the non-parametric LayerNorm used in Groeneveld et al. [65]. **(3) Performance improvements:** Besides some of the stability improvements which also impact performance, we also reduce the AdamW epsilon to 1.0E-08 from the 1.0E-05 used in Groeneveld et al. [65] to speed up convergence. Finally, we train OLMoE-1B-7B for significantly longer than all prior OLMo models amounting to 5T tokens and thus more than one epoch (1.3) following Muennighoff et al. [120]. We shuffle the pretraining dataset before starting the second epoch. To all data sources (Table C1), we apply a filter that removes all documents with a sequence of 32 or more repeated n-grams, where an n-gram is any span of 1 to 13 tokens. For the StarCoder subset, we also remove any document from a repository with fewer than 2 stars on GitHub, whose most frequent word constitutes over 30% of the document, or whose top-2 most frequent words constitute over 50% of the document. We shuffle all samples randomly at the beginning of each epoch and train for a total of 5.133T tokens. During

our annealing phase (final 100B tokens), we reshuffle the entire dataset and then linearly decay the learning rate from $5.0E-04$ to 0, following prior work [65, 89].

Adaptation For finetuning we use Open Instruct [186, 76].⁷ We filter all SFT samples to a length of fewer than 4096 tokens to match the sequence length of the model. Following Muennighoff et al. [121], we aggregate loss at the token level during SFT to improve performance on long generative tasks, such as AlpacaEval. We finetune in BF16 with a global batch size of 128 (4 H100 nodes with 8 GPUs each, a per device batch size of 2, and 2 gradient accumulation steps). We train for 2 epochs with a constant learning rate of $2.0E-5$. For DPO [136], we reduce the global batch size to 32 (4 H100 nodes with 8 GPUs each and a per device batch size of 1). We train for 3 epochs with a learning rate of $5.0E-7$ and a DPO beta of 0.1. Our adapted models are built on top of our annealed checkpoint, and we include the load balancing loss during both SFT and DPO based on our experiments in §B.3. Our preference tuning recipe is heavily optimized for DPO based on extensive experiments by Iverson et al. [76], thus for KTO [54] we experiment with a few settings in §G. Our final KTO adaptation uses the same hyperparameters as DPO, except that we use the RMSProp optimizer instead of Adam, which we use for SFT and DPO, and that we reduce the training duration to 1.3 epochs (5,000 steps) for KTO instead of the 3 epochs used for DPO.

Hardware We pretrain **OLMOE-1B-7B** on 256 H100 GPUs for approximately 10 days with NV-link interconnect across GPUs and InfiniBand interconnect across nodes. We also use H100 GPUs for all our experiments but some use a cluster with GCP TCPx interconnect across nodes instead. For adaptation, we use 32 H100 GPUs for 33 hours to instruction tune and for another 14 hours to preference tune via DPO. For KTO adaptation we use 8 H100 GPUs for 30 hours instead.

⁷Code: <https://github.com/allenai/open-instruct>

Table C3: **Pretraining hyperparameters of OLMoE-1B-7B and comparable models trained from scratch.** We highlight rows where **OLMoE-1B-7B** differs from OLMo-1B. Active params include vocab params. “?” = undisclosed settings, FFN = feed-forward network, Attn = Attention, LR = learning rate, WSD = Weight-Stable-Decay [74], LBL = load balancing loss, Inv Sq Root = Inverse Square Root decay [153], trunc = truncation, std = standard deviation, “varies” = stds that are layer or weight-dependent.

	OLMoE-1B-7B	JetMoE	OpenMoE	OLMo-1B (0724)
Dimension	2,048	2,048	2,048	2,048
Activation	SwiGLU	SwiGLU	SwiGLU	SwiGLU
FFN dimension	1,024	5,632	8,192	8,192
Vocab size	50,304	32,000	256,384	50,304
Attn heads	16	16	24	16
Num layers	16	24	32	16
Layer norm type	RMSNorm	RMSNorm	RMSNorm	non-parametric
Layer norm eps	1.0E-05	1.0E-05	1.0E-06	1.0E-05
QK-Norm	yes	no	no	no
Pos emb.	RoPE	RoPE	RoPE	RoPE
RoPE θ	10,000	10,000	10,000	10,000
Attention variant	full	MoA	full	full
Biases	-	MLP & Attn	-	-
Weight tying	no	yes	no	no
Init dist	trunc normal	?	?	normal
Init std	0.02	0.02	varies	varies
Init trunc	3×std	-	-	-
MoE layers	Every	Every	Every 6th	-
MoE layer type	dMoE	dMoE	ST-MoE	-
# Experts	64	8	32	1
# Activated	8	2	2	1
# Vocab params	103M	66M	525M	103M
# Active params	1.3B	2.2B	2.6B	1.3B
# Total params	6.9B	8.5B	8.7B	1.3B
Sequence length	4,096	4,096	2,048	4,096
Batch size (samples)	1,024	1,024	2,048	512
Batch size (tokens)	~4M	~4M	~4M	~2M
warmup steps	2,500	2,500	10,000	2,000
peak LR	4.0E-04	5.0E-04	0.01	4.0E-04
minimum LR	4.0E-05	5.0E-05	-	4.0E-05
optimizer	AdamW	AdamW	Adafactor	AdamW
weight decay	0.1	0.1	0.0	0.1
beta1	0.9	?	0.9	0.9
beta2	0.95	?	-	0.95
AdamW epsilon	1.0E-08	?	-	1.0E-05
LR schedule	cosine	WSD	Inv Sq Root	cosine
gradient clipping	global 1.0	global 1.0	global 1.0	global 1.0
gradient reduce dtype	FP32	?	?	FP32
optimizer state dtype	FP32	?	?	FP32
LBL weight	0.01	0.01	0.01	-
Router z-loss weight	0.001	0.001	0.0001	-
Pretraining tokens	5,033B	1,000B	1,100B	2,000B
Annealing tokens	100B	250B	-	50B
Annealing schedule	linear	-	-	linear
Annealing min LR	0	-	-	0

D Evaluation setup

Table D1: **Summary of downstream evaluation during and after pretraining (OLMES)**. ARC-C and ARC-E refer to ARC-Challenge and -Easy, CSQA=CommonsenseQA, OBQA=OpenBookQA, CF=Completion/Cloze formulation, MCF=Multiple-choice formulation, pmi=pointwise-mutual-information, Var=variants referring to the use of few-shots varying from 0-5.

Dataset (↓)	During pretraining				After pretraining (OLMES [67])			
	Format	Shot	Norm	Split	Format	Shot	CF Norm	Split
ARC-C [34]	CF	0	token	val	max(MCF,CF)	5	pmi	test
ARC-E [34]	CF	0	none	val	max(MCF,CF)	5	character	test
BoolQ [33]	CF	0	none	val	max(MCF,CF)	5	none	val
COPA [63]	CF	0	none	val	-	-	-	-
CSQA [168]	CF	0	token	val	max(MCF,CF)	5	pmi	val
HellaSwag [205]	CF	0	token	val	max(MCF,CF)	5	character	val
MMLU [70]	MCF	5	none	val	max(MCF,CF)	5	character	test
MMLU Var	CF	0-5	token	val	-	-	-	-
OBQA [116]	CF	0	token	val	max(MCF,CF)	5	pmi	test
PIQA [20]	CF	0	token	val	max(MCF,CF)	5	character	val
SciQ [190]	CF	0	none	val	-	-	-	-
SocialQA [148]	CF	0	token	val	max(MCF,CF)	5	character	val
Winogrande [146]	CF	0	none	val	max(MCF,CF)	5	none	val

During pretraining We evaluate using a similar in-loop evaluation setup as Groeneveld et al. [65], with the addition of more tasks such as CommonsenseQA, PIQA, and different implementations of MMLU. Following Groeneveld et al. [65], for the majority of the tasks, we perform 0-shot evaluation using the Completion/Cloze formulation (CF), ranking each answer string using language model probabilities. In terms of probability normalization, there is either no normalization (none) or normalization by the number of tokens in the answer (token) when ranking solely based on probability may heavily favor shorter answers [24]. For MMLU, the in-loop evaluation also includes a setup where we increase the total number of instances by including a range of 0-shot to 5-shot setups together as we found this provides smoother trends as the training proceeds (“MMLU Var”). We also include the Multiple-choice formulation (MCF) version of MMLU, scoring prediction of answer labels like A/B/C/D, which generally starts to rise only later in training as models only gain the multiple-choice capability later (at around 1T tokens for **OLMOE-1B-7B** in Figure F3). We also evaluate perplexity on selected validation sets from Paloma [110, 142, 59, 161, 95, 114]. All code used for evaluation during pretraining is at <https://github.com/allenai/OLMo/tree/61ac104d616ec5435db225796e5c7532c9abd95a/olmo/eval>.

After pretraining - OLMES We perform evaluations following the OLMES evaluation standard [67], with the suite of tasks in the original paper. OLMES (Open Language Model Evaluation Standard) is a standard for reproducible LM evaluations that is open, practical, and documented, providing recommendations guided by experiments and results from the literature [19, 60, 64]. It is designed to support comparisons between smaller base models that require the Cloze formulation of multiple-choice questions against larger models that can utilize the Multiple-choice formulation. To make our evaluations reproducible, we follow OLMES in prompt formatting, choice of in-context examples, probability normalization, task formulation, as well as all other details. We summarize this setup in Table 2 and refer to Gu et al. [67] for more details.

After pretraining - DCLM For results on the DCLM tasks [89] in Table F2, we precisely follow their setup using the evaluation code released by the authors at <https://github.com/mlfoundations/dclm>. “Core” results are the low variance tasks in their evaluation code, while “Extended” corresponds to the heavy tasks.





After adaptation After supervised finetuning and direct preference optimization, we evaluate models using a subset of the evaluations and the same overall setup used in Ivison et al. [76] and Wang et al. [186]. We cover a wide range of model capabilities in our evaluation suite including

coding (HumanEval [28]), general and mathematical reasoning (Big Bench Hard [167], GSM8k [35]), world knowledge (MMLU), general instruction following (AlpacaEval 1.0 [92], not the length-controlled variant [51]), precise instruction following (IFEval [215]) and safety (XSTest [145]). We refer to Wang et al. [186] for more details on each benchmark.





E Openness of models

We list the openness of various models summarized in Figure 1. We exclude Switch Transformers [56], as it was published over three years ago and is very different from more recent MoE models (MLM objective, Encoder-decoder, etc.).





Grok-86B-314B [194]

-  **Model:** Their model is licensed under the open-source Apache 2.0 license.
-  **Data:** Unavailable.
-  **Code:** Unavailable.
-  **Logs:** Unavailable.





Mixtral-39B-141B and Mixtral-13B-42B [79]

-  **Model:** Their model is licensed under the open-source Apache 2.0 license.
-  **Data:** Unavailable.
-  **Code:** Unavailable.
-  **Logs:** Unavailable.





DBRX-36B-132B [40]

-  **Model:** The model is licensed under a custom non-open-source license⁸ with additional use-case restrictions.⁹
-  **Data:** Unavailable.
-  **Code:** They use closed-source custom adaptations of their public libraries LLM-foundry, composer, and megablocks.¹⁰
-  **Logs:** Unavailable.

Skywork-MoE-22B-146B [189]

-  **Model:** The model is licensed under a custom non-open-source license.¹¹
-  **Data:** Unavailable.
-  **Code:** Unavailable.
-  **Logs:** Unavailable.

DeepSeekV2-21B-236B [43] and DeepSeekMoE-3B-14B [39]

-  **Model:** The models are licensed under custom non-open-source licenses.¹²
-  **Data:** Unavailable.
-  **Code:** Unavailable.
-  **Logs:** Unavailable.

⁸<https://www.databricks.com/legal/open-model-license>





⁹<https://www.databricks.com/legal/acceptable-use-policy-open-model>

¹⁰<https://github.com/databricks/dbrx>





¹¹<https://github.com/SkyworkAI/Skywork/blob/main/Skywork%20Community%20License.pdf>

¹²<https://github.com/deepseek-ai/DeepSeek-MoE/blob/main/LICENSE-MODEL> and <https://github.com/deepseek-ai/DeepSeek-V2/blob/main/LICENSE-MODEL>





Arctic-17B-480B [160]

-  **Model:** The model is licensed under the open-source Apache 2.0 license.
-  **Data:** They describe their mixture but do not release it.¹³
-  **Code:** Unavailable.
-  **Logs:** Unavailable.





Qwen2-14B-57B [178]

-  **Model:** The model is licensed under the open-source Apache 2.0 license.
-  **Data:** Unavailable.
-  **Code:** Unavailable.
-  **Logs:** Unavailable.





Jamba-12B-52B [96]

-  **Model:** The model is licensed under the open-source Apache 2.0 license.
-  **Data:** Unavailable.
-  **Code:** Unavailable.
-  **Logs:** Unavailable.





Qwen1.5-3B-14B [178]

-  **Model:** The model is licensed under a custom non-open-source license.¹⁴
-  **Data:** Unavailable.
-  **Code:** Unavailable.
-  **Logs:** Unavailable.

JetMoE-2B-9B [156]

-  **Model:** The model is licensed under the open-source Apache 2.0 license.
-  **Data:** They describe their mixture but do not release it.
-  **Code:** They make their fork of megablocks publicly available,¹⁵ however, their Megatron-LM training code is not available.¹⁶
-  **Logs:** Unavailable.

OpenMoE-2B-9B [197]

-  **Model:** The model is licensed under the open-source Apache 2.0 license.
-  **Data:** They make scripts for recreating their data available.
-  **Code:** They make their code available.¹⁷
-  **Logs:** Unavailable.

¹³<https://medium.com/snowflake/snowflake-arctic-cookbook-series-arctics-approach-to-data-b81a8a0958bd>

¹⁴<https://hf.co/Qwen/Qwen1.5-MoE-A2.7B/blob/main/LICENSE>

¹⁵<https://github.com/yikangshen/megablocks>

¹⁶<https://hf.co/jetmoe/jetmoe-8b/discussions/5#661ee52c03251697a0b155cc>

¹⁷<https://github.com/XueFuzhao/OpenMoE/tree/main?tab=readme-ov-file#training-with-tpugpu>

OLMoE-1B-7B

- **Model:** The model is licensed under the open-source Apache 2.0 license.
- **Data:** The data is licensed under the open-source ODC-By 1.0 license.
- **Code:** The code is licensed under the open-source Apache 2.0 license.
- **Logs:** Logs are available with the same open-source license as the code (Apache 2.0).

F Additional evaluation

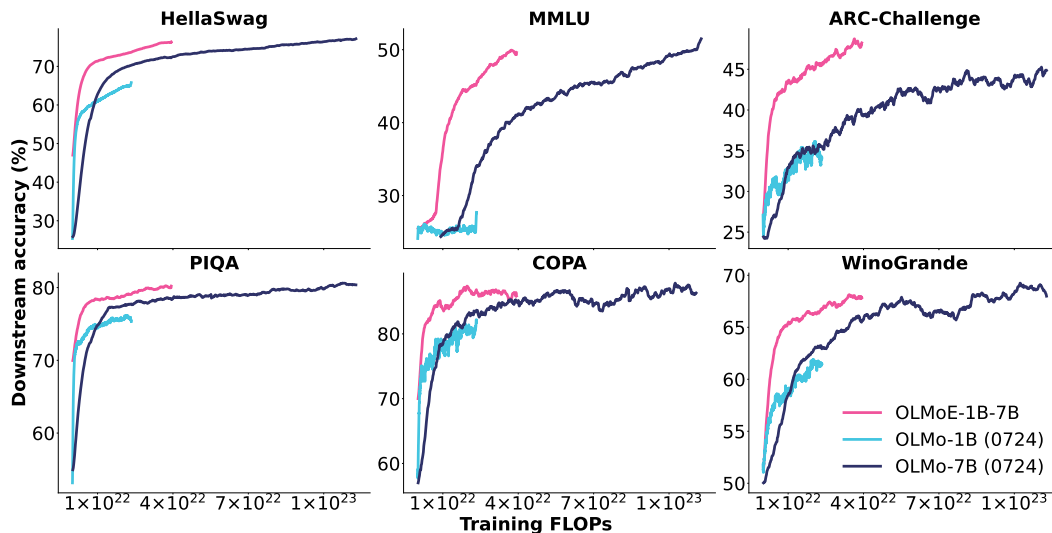


Figure F1: **Evaluation of OLMoE-1B-7B and the current best OLMo models during pretraining.** **OLMoE-1B-7B** differs from the OLMo models in its MoE architecture, several training hyperparameters, and its training dataset, see §2. A version of this plot with tokens as the x-axis and markers where annealing starts is in §F. More results, logs, and configurations: <https://wandb.ai/ai2-11m/olmo/reports/Plot-OLMoE-1B-7B-vs-OLMo-7B-vs-OLMo-1B--Vml1dzo40TcyMjEz>

During pretraining In Figure F1 we benchmark the performance of **OLMoE-1B-7B** during pretraining with the current best OLMo models [65] on commonly used downstream tasks. We find that across all tasks **OLMoE-1B-7B** reaches better performance with less compute (FLOPs) than the dense OLMo models. **OLMoE-1B-7B** matches or outperforms OLMo-7B at the end of training despite **OLMoE-1B-7B** having used less than half as many FLOPs for training and using only 1B active parameters. This is likely a result of the dataset and modeling changes we make to the OLMo setup including MoE-related changes, stability, and performance improvements, outlined in §C. §F contains training and validation loss plots showing very smooth loss curves without major loss spikes during the 5T tokens of our pretraining.

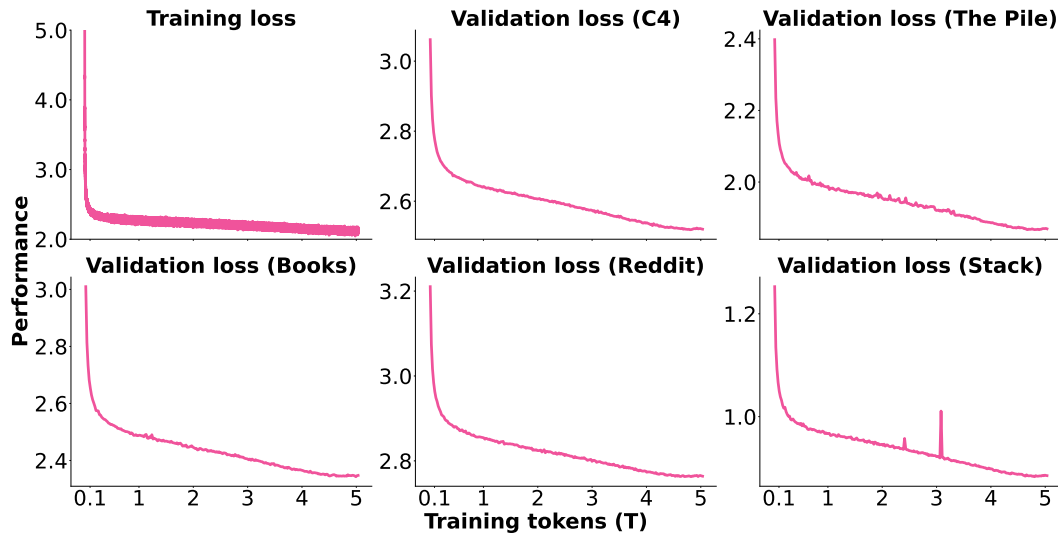


Figure F2: **Losses of OLMoE-1B-7B during training.** The Books, Reddit, and Stack [84] datasets are from Dolma 1.7 [161] via Paloma [110]. More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmo/reports/Plot-OLMoE-1B-7B--Vml1dzo40TcyMjU3>




Figure F3: **Evaluation of OLMoE-1B-7B and the current best OLMo models during pretraining.** Grey vertical lines correspond to where the respective run enters anealing with the 1st line being for OLMo-7B, the 2nd for OLMo-1B, and the third for OLMoE-1B-7B. Figure F1 is a version of this plot with training FLOPs as the x-axis. More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmo/reports/Plot-OLMoE-1B-7B-vs-OLMo-7B-vs-OLMo-1B--Vml1dzo40TcyMjEz>

Table F1: More results on OLMES. † indicates use of the MCF score, see §D. See Table 2 for details on naming and a summary of these results.

Model	ARC.C	ARC.E	BoolQ	CSQA	HSwag	MMLU	OBQA	PIQA	SIQA	WinoG	Avg
LMs with ~7-9B active parameters											
Mistral-7B	78.6 [†]	90.8 [†]	89.3	72.4 [†]	83.0	64.0 [†]	80.6 [†]	82.8	71.3 [†]	77.9	79.1
OLMo-7B (0724)	68.0 [†]	85.7 [†]	85.3	85.4 [†]	80.5	54.9 [†]	67.6 [†]	79.3	76.1 [†]	73.2	75.6
DCLM-7B	79.8 [†]	92.3 [†]	87.0	77.0	82.3	64.4 [†]	79.6 [†]	80.1	71.2 [†]	77.3	79.1
Llama2-7B	54.2	84.0	86.1	74.2	78.9	46.2 [†]	57.8	77.5	59.6	71.7	69.0
Llama3.1-8B	79.5 [†]	91.7 [†]	88.5	74.3 [†]	81.6	66.9 [†]	78.6 [†]	81.1	71.4 [†]	76.6	79.0
Gemma2-9B	89.5 [†]	95.5 [†]	89.4	78.8 [†]	87.3 [†]	70.6 [†]	88.4 [†]	86.1 [†]	76.0 [†]	78.8	84.0
LMs with ~2-3B active parameters											
StableLM-2B	50.6 [†]	75.3	82.3	70.4 [†]	70.3	40.4 [†]	56.6 [†]	75.6	64.3 [†]	65.8	65.1
Gemma2-3B	67.5 [†]	84.3 [†]	83.6	66.4 [†]	74.6	53.3 [†]	68.8 [†]	78.5	64.7 [†]	71.8	71.4
JetMoE-2B-9B	61.4 [†]	81.9 [†]	85.7	75.3 [†]	81.7	49.1 [†]	68.0 [†]	80.3	71.3 [†]	70.7	72.5
OpenMoE-3B-9B	29.3	50.6	63.2	21.5	44.4	27.4	34.6	63.3	42.9	51.9 [†]	42.9
DeepSeek-3B-16B	53.4	82.7	81.9	72.7	80.4	45.5 [†]	58.4	80.1	59.9	73.2	68.8
Llama3.2-3B	69.6 [†]	85.1 [†]	78.3	69.0	77.0	57.8 [†]	67.2 [†]	77.4	64.9 [†]	69.9	71.6
Qwen1.5-3B-14B	77.4 [†]	91.6 [†]	85.0	81.4 [†]	80.0	62.4 [†]	80.6 [†]	81.0	74.1 [†]	72.3	78.6
LMs with ~1B active parameters											
OLMo-1B (0724)	36.4	53.5	66.8	42.4	67.5	32.1	44.2	74.0	45.2	62.9	52.5
TinyLlama-1B	38.1	69.5	63.6	61.1	60.8	33.6	45.0	71.7	50.4	60.1	55.4
Pythia-1B	31.4	63.4	56.8 [†]	50.9	48.0	31.1	40.4	68.9	46.4	52.7	49.0
Llama3.2-1B	43.5	71.6	69.4	59.6	67.3	38.2	42.0	73.7	52.0	62.5	58.0
Zamba2-1B	55.0 [†]	85.4	76.1	70.1	73.4	44.73 [†]	59.8 [†]	76.6	58.4	67.2	66.7
DCLM-1B	57.6 [†]	79.5	80.9	71.3	75.1	48.5 [†]	60.0 [†]	76.6	60.5 [†]	68.1	67.8
OLMoE-1B-7B	62.1[†]	84.2	79.2	72.9	80.0	54.1[†]	65.4[†]	79.8	63.0[†]	70.2	71.1

Table F2: **DCLM evaluation metrics on the Core and Extended task subsets [89]**. * = Core tasks. “annealed” is the final pretraining checkpoint we use for **OLMOE-1B-7B** and was annealed from the checkpoint at step 1,200,000. We left the non-annealing pretraining run train a little longer resulting in the 1,220,000 checkpoint.

OLMOE-1B-7B checkpoint (→)	step 1,200,000	step 1,220,000	annealed	OLMo-1B	OLMo-7B
AGI Eval LSAT-AR*	24.3	26.5	28.7	28.3	28.3
AGI Eval LSAT-LR	40.2	38.6	37.3	30.2	42.9
AGI Eval LSAT-RC	47.4	43.7	46.6	23.5	61.6
AGI Eval SAT-En	55.3	54.9	52.9	28.2	73.8
AGI Eval SAT-Math CoT	5.5	4.1	6.4	1.8	6.8
AQuA CoT	2.4	2.9	2.0	2.9	6.1
ARC Challenge*	53.3	53.4	53.8	34.6	48.1
ARC Easy*	77.1	78.5	77.7	64.4	75.9
BBQ	49.8	48.3	50.6	45.8	67.2
BigBench CS Algorithms*	47.1	50.2	47.2	47.5	53.6
BigBench Conceptual Combinations	51.5	50.5	56.3	31.1	68.0
BigBench Conlang Translation	3.7	6.1	7.3	4.3	7.3
BigBench Dyck Languages*	19.3	15.9	21.5	26.6	22.2
BigBench Elementary Math QA	26.2	27.0	26.9	26.2	30.4
BigBench Language Identification*	31.9	34.0	31.0	27.0	39.1
BigBench Logical Deduction	26.6	25.3	24.6	23.6	27.3
BigBench Misconceptions	59.8	55.3	62.6	55.7	58.0
BigBench Novel Concepts	62.5	62.5	65.6	43.8	53.1
BigBench Operators*	36.2	34.3	33.8	23.8	45.2
BigBench QA Wikidata*	68.2	68.8	69.2	67.0	69.9
BigBench Repeat Copy Logic*	15.6	15.6	18.8	3.1	9.4
BigBench Strange Stories	66.7	68.4	69.5	53.4	66.1
BigBench Strategy QA	56.2	58.1	57.0	51.5	68.6
BigBench Understanding Fables	47.1	44.4	47.6	28.0	61.4
BoolQ*	73.3	72.8	73.2	63.7	83.9
COPA*	81.0	80.0	78.0	75.0	77.0
CoQA*	43.7	44.4	43.7	3.4	45.4
CommonsenseQA*	67.2	67.0	69.3	19.6	86.0
Enterprise PII Classification	52.3	53.7	52.2	57.3	50.6
GPQA Diamond	22.2	21.2	19.7	19.7	20.2
GPQA Main	24.8	22.3	22.5	20.3	23.0
GSM8K CoT	6.4	7.4	7.4	4.9	30.6
HellaSwag 0-shot*	76.0	76.0	77.0	65.8	76.7
HellaSwag 10-shot*	77.6	77.5	78.6	66.3	78.9
Jeopardy*	48.8	48.7	50.3	22.6	46.5
LAMBADA*	72.7	72.2	73.3	61.1	71.8
LogiQA	34.9	34.3	34.6	28.7	31.0
MMLU Few-shot	52.2	51.9	53.3	28.4	55.1
MMLU Zero-shot	41.6	42.7	43.3	26.2	50.0
Math QA	26.4	27.1	27.5	24.1	29.8
OpenBookQA*	41.4	44.0	44.8	36.6	43.4
PIQA*	81.3	81.2	82.0	76.4	81.7
PubMedQA	56.1	46.6	57.9	0.2	57.9
SQuAD*	52.9	52.4	52.4	0.0	65.5
SVAMP CoT	30.0	28.0	33.0	14.3	44.7
Simple Arithmetic, no spaces	17.6	18.1	20.1	1.2	15.3
Simple Arithmetic, with spaces	19.5	20.6	22.1	1.8	16.0
Social IQA	71.5	70.7	69.3	69.5	84.4
Trivia QA	54.2	53.0	55.9	25.1	51.8
Winogender Female	50.0	46.7	50.0	41.7	58.3
Winogender Male	55.0	58.3	60.0	63.3	58.3
Winograd*	82.8	83.2	84.6	79.9	83.2
Winogrande*	68.0	68.5	69.0	61.8	67.6
Core	46.3	46.5	47.2	30.2	49.8
Extended	31.3	30.9	32.5	16.9	37.0

G Additional experiments




Figure G1: **Adding Reddit or FLAN to OLMoE-Mix.** More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-Adding-Reddit-FLAN--Vml1dzo40Tg1NTg4>

Adding Reddit or FLAN to OLMoE-Mix In Figure G1 we benchmark adding the Reddit or FLAN [188] subsets of Dolma 1.7 [161] to our pretraining data mix (§2). Overall, we do not find either one to lead to consistent gains, thus we do not use them in our final data mix.

Load balancing precision Fedus et al. [56] selectively perform operations related to routing in full precision (FP32) to improve stability. In Figure G2, we test whether computing the load balancing loss in full precision improves stability, but do not find it to reduce spikes. Thus, we stick with bfloat16 (BF16).

Noise upcycling For the creation of Qwen2-MoE [199, 178, 13], the authors add 50% of gaussian noise to feedforward networks before continuing training in an upcycled setup [85]. Komatsuzaki et al. [85] also report that they experimented with adding noise but did not find it beneficial. In Figure G3, we experiment with regular upcycling versus adding noise by randomly replacing 50% of each MLP with numbers drawn from a normal distribution with a standard deviation of 0.02




Figure G2: **Load balancing precision.** More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-FP32-LBL--Vmlldzo4NDMxNDA4>

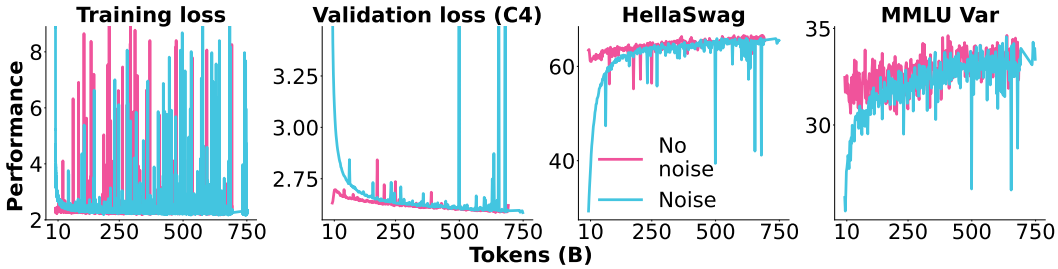


Figure G3: **Adding noise to the upcycled checkpoint.** More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-Noise-upcycle---Vmlldzo4NDA3MzI2>

following. We find that after 700 billion tokens, the no noise variant still performs slightly better but both appear to converge to the same performance. If training further, it is possible that the noise variant eventually outperforms the no noise variant, but at that point, it may make more sense to just train the MoE from scratch (§B.1.5).




Figure G4: **Sharing the same MoE across layers versus a regular dense LM.** The number of experts in the MoE is equivalent to its number of layers. Thus, because the MoE is shared across layers, it has the same number of total and active parameters as the dense model. More results, logs, and configurations: <https://wandb.ai/ai2-llm/olmoe/reports/Plot-Shared-vs-Dense--Vmlldzo4NDI0MTc5>

Shared Layer Some work has investigated Mixture-of-Experts with weights shared across layers in the context of Universal Transformers [169, 37, 45]. We test whether layer-shared Mixture-of-Experts can beat non-shared dense models in Figure G4. The layer-shared MoE uses a load balancing loss that is applied at the model level rather than at the layer level. This gives the model more flexibility by allowing it to completely deactivate certain experts for some layers and even emulate a dense model by always activating one separate expert for each layer. This makes it a generalization of the dense model which motivated our hypothesis that it may perform better than the dense model. However, in practice, we find that both perform similarly with the regular dense models even maintaining a small advantage on validation loss and HellaSwag. One possible advantage of layer-shared MoEs

is that they can allow for better load balancing at inference. If prompts come in continuously, then newly incoming prompts can be batched with previous prompts that have already passed through several layers and sent through the MoE module together, as the MoE module is the same regardless of whether it is the first or last layer. Sharing also reduces throughput by around 20% during training, which further motivates our decision not to use it for **OLMoE-1B-7B**.

KTO experiments In Table G1 we experiment with the number of steps (5,000 vs. 10,000) and the optimizer (Adam [83] vs. RMS) used for KTO [54]. Based on these experiments we use the RMS optimizer and the checkpoint at 5,000 steps in §B.3.

Table G1: **KTO adaptation experiments.** 5,000 and 10,000 steps correspond to 1.3 and 2.6 epochs on our adaptation dataset (§2), respectively.

Task (→) Setup (→) Metric (→)	MMLU	GSM8k	BBH	Human- Eval	Alpaca- Eval 1.0	XSTest	IFEval	Avg
	0-shot	8-shot CoT	0-shot	0-shot	0-shot	0-shot	0-shot	0-shot
	EM	EM	EM	Pass@10	%win	F1	Loose Acc	0-shot
KTO, 5,000 steps, RMS	51.2	45.5	34.1	57.1	81.6	86.6	47.5	57.7
KTO, 10,000 steps, RMS	51.0	41.0	34.7	53.8	81.0	62.3	47.5	54.2
KTO, 5,000 steps, Adam	51.2	42.0	35.3	55.6	81.0	84.5	46.6	56.0
KTO, 10,000 steps, Adam	51.0	43.0	34.1	54.9	79.7	62.7	47.5	53.3

H Analysis

H.1 Details of analysis in §4

Router saturation We define router saturation as the proportion of expert activations at some intermediary checkpoint at time t that matches the expert IDs activated at some final checkpoint over the same dataset:

$$\text{Router Saturation}(t) = \frac{1}{N} \sum_{i=1}^N \frac{|\mathcal{E}_i^{(t)} \cap \mathcal{E}_i^{(T)}|}{k}, \quad (5)$$

where:

- N : The total number of tokens in the dataset.
- k : The number of top- k experts activated per input token. While we train with $k = 8$ (§2), we also analyze $k = 1$ by only looking at the expert with the highest routing probability.
- $\mathcal{E}_i^{(t)}$: The set of k experts activated for the i th token at the t th checkpoint.
- $\mathcal{E}_i^{(T)}$: The set of k experts activated for the i th token at the final checkpoint T .
- $|\mathcal{E}_i^{(t)} \cap \mathcal{E}_i^{(T)}|$: The number of common experts activated for the i th token between the t th and final checkpoints.

Router saturation thus corresponds to whether the router weights are still learning which expert will process certain data. A value of 100% indicates that the router at the intermediate checkpoint will route to the same experts as the final checkpoint router. However, even at 100% saturation the router weight can still change and adapt the exact router probability for each expert. These probabilities are used to scale the output of the respective expert in the model. For **OLMoE-1B-7B** with its 64 experts, random routing equals a saturation of $1/64 = 1.6\%$ for $k = 1$ and $8/64 = 12.5\%$ for $k = 8$.

Expert co-activation We define expert co-activation as the proportion of times two specific experts, E_i and E_j , are simultaneously activated out of the total number of activations of one of those experts:

$$\text{Expert co-activation}(E_i, E_j) = \frac{N_{E_i, E_j}}{N_{E_i}}, \quad (6)$$

where:

- E_i : The first expert.
- E_j : The second expert.
- N_{E_i, E_j} : The number of times experts E_i and E_j are activated together.
- N_{E_i} : The total number of times expert E_i is activated.

A co-activation of 100% indicates that if E_i is activated, E_j is also always activated. A value of 0% indicates that the experts never co-occur. If multiple expert pairs have high co-activation, it may suggest that these experts could be merged, benefiting less from keeping them separate. In a distributed setup, we could place highly co-activated experts on the same device to reduce communication costs during model inference.

Domain specialization We define domain specialization as the proportion of tokens from a particular domain D that get routed to a particular expert E_i :

$$\text{Domain specialization}(E_i, D) = \frac{N_{E_i, D}^{(k)}}{N_D}, \quad (7)$$

where:

- E_i : The i th expert in the model.
- D : The domain from which the data originates.
- k : The number of experts considered (e.g., $k = 8$ means considering the top 8 experts with the highest routing probabilities).
- $N_{E_i, D}^{(k)}$: The number of tokens from domain D for which E_i is among the top- k selected experts.
- N_D : The total number of tokens from domain D processed by the MoE.

Domain specialization thus refers to the specialization of expert E_i to domain D . A value of 100% indicates that all data from that domain is routed to E_i , whereas 0% indicates the expert is never used for that domain and can be removed from the model without affecting performance in that domain.

Vocabulary specialization We define vocabulary specialization as the proportion of tokens with a token ID x (also called vocabulary element) that are routed to one particular expert E_i out of all experts in that layer:

$$\text{Vocabulary specialization}(E_i, x) = \frac{N_{x, E_i}^{(k)}}{N_x}, \quad (8)$$

where:

- E_i : The i th expert in the model.
- x : The token ID being analyzed.
- k : The number of experts considered (e.g., $k = 8$ means considering the top 8 experts with the highest routing probabilities).
- N_{x, E_i} : The number of times input data is routed to E_i for x .
- N_x : The total number of times input data is routed across all experts for x .

Vocabulary specialization thus refers to how specialized a particular expert is on some vocabulary item. We distinguish input and output variants of this specialization, where x is either the input token ID or the next output token ID (either the ground-truth next token ID or the token ID predicted by the model). A value of 100% indicates that for all occurrences of that vocabulary element, input data is routed to E_i , whereas 0% indicates an expert that is fully irrelevant for that vocabulary element and can be effectively removed from the model without affecting performance whenever the token ID appears.

H.2 Additional Analysis




Figure H1: **Co-activation among experts of OLMoE-1B-7B on a random 0.5% of the C4 validation data.** We display the 32 experts with the highest maximum co-activation score via their expert IDs on the x- and y-axis. See Figure 3 for layer 7.

Table H1: **Vocabulary specialization in the 7th layer of OLMoE-1B-7B.** We use $k = 1$ (Equation 8) and a random 0.5% of the C4 validation data excluding token IDs with < 10 appearances. See Table 4 for more.

Expert ID	Input token IDs	Predicted output token IDs
58	“ (100%) ” (100%) ‘ (94%) ’ (92%) “ (92%) ((92%) ” (90%) ’ (89%) “ (88%) \$ (87%) [(87%) £ (86%)	such (100%) 486 (100%) see (95%) which (91%) driving (91%) UK (90%) who (88%) including (88%) normal (88%)
7	Him (100%) inde (100%) Jesus (98%) God (90%) pray (81%) Holy (80%) Quran (80%) God (77%) Lord (76%) glory (75%) Spirit (66%) Christ (65%)	rella (100%) Him (94%) sin (90%) prince (80%) glory (72%) Jesus (69%) Lord (68%) Christ (65%) Spirit (55%) Holy (53%) God (50%) Prayer (50%)
37	Sunday (100%) Tuesday (100%) Thursday (100%) Olympic (100%) Christmas (100%) rugby (100%) Championship (100%) weekends (100%)	days (91%) anniversary (90%) month (88%) week (84%) mpi (83%) semester (81%) mand (80%) Olympics (78%) cent (76%) season (76%) perm (75%)
0	ESM (100%) icillin (100%) agra (98%) aust (96%) asa (93%) pills (92%) (85%) uk (82%) login (82%) doc (81%) generic (81%) cd (81%) Essay (81%) password (81%) Content (80%)	*, (100%) sil (96%) pills (91%) vi mg (90%) xen (87%) pharmacy (87%) gener (85%) aust (82%) mg (75%) Content (75%) uk (73%) THAT (73%) dispens (68%) icillin (68%) generic (66%)
3	grandmother (92%) brother (91%) Daisy (83%) daughter (78%) mum (75%) father (72%) wife (70%) husband (70%) lady (63%) dad (62%) boy (61%)	hood (36%) mother (35%) inde (31%) boy (29%) girl (28%) married (27%) tri (21%) Gab (20%) died (18%) taught (14%) lived (13%) knew (10%)




Figure H2: Vocabulary specialization for OLMoE-1B-7B when considering all 8 activated experts. Equivalent to $k = 8$ in Equation 8.




Figure H3: Vocabulary specialization for Mixtral-8x7B when considering all 2 activated experts. Equivalent to $k = 2$ in Equation 8.




Figure H4: **Vocabulary specialization across domains of OLMoE-1B-7B (top) and Mixtral-8x7B (bottom).** We visualize how often token IDs get routed to specific experts. We only include IDs that appear at least 8 times in the various corpora. Vertical gray lines correspond to uniform routing ($8/64=12.5\%$ for OLMoE-1B-7B as it has 64 experts, 8 of which are activated; $2/8=25\%$ for Mixtral as it has 8 experts, 2 of which are activated). For example, among all token IDs in GitHub that get routed to Expert 0 at least 8 times for OLMoE-1B-7B, $\sim 40\%$ of them get routed to Expert 0 with a probability of $\sim 100\%$ (upper left) indicating that Expert 0 is specialized on those token IDs. For OLMoE-1B-7B there is much frequency at the routing probability extremes (0% or 100%) indicating that these experts exclusively focus on certain token IDs, especially for *specific domains* (§4.3) like GitHub and arXiv.




Figure H5: **Load imbalances in selective layers after adaptation.** We visualize how often tokens from our instruction tuning dataset (§2) get routed to the 8 active experts out of the 64 total experts ($k = 1$ in Equation 7). Horizontal gray lines correspond to uniform routing ($8/64=12.5\%$ per expert). Although we run SFT and DPO without loss balancing loss (§B.3), we observe that the load distribution does not change substantially.




Figure H6: **Domain specialization of OLMoE-1B-7B (top) vs. Mixtral-8x7B (bottom) of the top-1 routed expert.** We visualize how often tokens from different domains get routed to the 64 (OLMoE) or 8 (Mixtral) experts at the end of pretraining. Unlike in Figure 4, here we only consider tokens routed to the top-1 expert ($k = 1$ in Equation 7). Horizontal gray lines correspond to uniform routing ($1/64=1.56\%$ per expert for OLMoE-1B-7B and $1/8=12.5\%$ for Mixtral).




Figure H7: **OLMoE-1B-7B token routing across layers.** We visualize how often tokens from different domains get routed to a pair of experts across layers under top-1 routing, corresponding to Figure H6. The size of each rectangle is proportional to the total number of tokens an expert receives, while the flow between two experts shows the proportion of tokens routed to both experts. We only show experts that receive tokens 50% above random chance and use stronger coloring for larger flows. We observe some instances of cross-layer coordination between pairs of experts, e.g., expert 27 in layer 7 and expert 57 in layer 15 process a substantial fraction of Wikipedia tokens together. The flows between layers $0 \rightarrow 7$ and $7 \rightarrow 15$ are independent in this visualization.




Figure H8: **Mixtral-8x7B token routing across layers.** We visualize how often tokens from different domains get routed to a pair of experts across layers under top-1 routing, corresponding to Figure H6. The size of each rectangle is proportional to the total number of tokens an expert receives, while the flow between two experts shows the proportion of tokens routed to both experts. The flows between layers 0 \rightarrow 7 and 7 \rightarrow 15 are independent in this visualization.

I Artifacts

Table I1: **All artifacts released and used in this work.** We point from the name used for a given artifact in this work (e.g. Figure 1) to the URL where it can be obtained.

Artifact	Public link
OLMoE-1B-7B	https://hf.co/allenai/OLMoE-1B-7B-0924
OLMoE-1B-7B-INSTRUCT	https://hf.co/allenai/OLMoE-1B-7B-0924-Instruct
OLMoE-1B-7B-SFT	https://hf.co/allenai/OLMoE-1B-7B-0924-SFT
OLMoE-Mix	https://hf.co/datasets/allenai/OLMoE-mix-0924
SFT data	https://hf.co/datasets/allenai/tulu-v3.1-mix-preview-4096-OLMoE
KTO/DPO data	https://hf.co/datasets/allenai/ultrafeedback_binarized_cleaned
Code	https://github.com/allenai/OLMoE
Logs	https://wandb.ai/ai2-llm/olmoe/reports/OLMoE-1B-7B-0924--Vmlldzo40TcyMjU3
BLOOM-7B	https://hf.co/bigscience/bloom-7b1
DeepSeekMoE-3B-16B	https://hf.co/deepseek-ai/deepseek-moe-16b-base
DeepSeekMoE-3B-16B+chat	https://hf.co/deepseek-ai/deepseek-moe-16b-chat
DCLM-1B	https://hf.co/TRI-ML/DCLM-1B
DCLM-7B	https://hf.co/TRI-ML/DCLM-7B
Falcon-7B	https://hf.co/tiiuae/falcon-7b
Gemma2-3B	https://hf.co/google/gemma-2-2b
Gemma2-9B	https://hf.co/google/gemma-2-9b
JetMoE-2B-9B	https://hf.co/jetmoe/jetmoe-8b
JetMoE-2B-9B+SFT	https://hf.co/jetmoe/jetmoe-8b-sft
JetMoE-2B-9B+Chat	https://hf.co/jetmoe/jetmoe-8b-chat
Llama-7B	https://hf.co/huggyllama/llama-7b
Llama2-7B	https://hf.co/meta-llama/Llama-2-7b-hf
Llama3.1-8B	https://hf.co/meta-llama/Meta-Llama-3.1-8B
MPT-7B	https://hf.co/mosaicml/mpt-7b
Mistral-7B	https://hf.co/mistralai/Mistral-7B-v0.1
Mixtral-8x7B	https://hf.co/mistralai/Mixtral-8x7B-v0.1
OLMo-1B (0724)	https://hf.co/allenai/OLMo-1B-0724-hf
OLMo-7B (0724)	https://hf.co/allenai/OLMo-7B-0724-hf
OpenMoE-3B-9B	https://hf.co/OrionZheng/openmoe-8b
Pythia-7B	https://hf.co/EleutherAI/pythia-6.9b
Qwen1.5-3B-14B	https://hf.co/Qwen/Qwen1.5-MoE-A2.7B
Qwen1.5-3B-14B+Chat	https://hf.co/Qwen/Qwen1.5-MoE-A2.7B-Chat
StableLM2-2B	https://hf.co/stabilityai/stablelm-2-1_6b
TinyLlama-1B	https://hf.co/TinyLlama/TinyLlama_v1.1

J Limitations and future work

We highlight four key limitations with this release of **OLMoE-1B-7B**. We look forward to addressing these issues in future iterations of **OLMoE**.

More parameters **OLMoE-1B-7B** has 7B total parameters out of which 1B are activated for each input token. This small size makes **OLMoE-1B-7B** very cheap to use, yet we demonstrate in this work that it outperforms much more expensive models (Figure 1). However, using only 1B parameters for each input token also limits the capabilities of **OLMoE-1B-7B** as seen by its performance compared to models that use $>7\times$ more parameters, such as Llama3.1-8B in §3. While it may be possible that more parameters are not needed to match 8B models and beyond [81], in the short-term adding parameters is an easy way to improve the performance of **OLMoE**, at least allowing the model to utilize more than 1B parameters per input, possibly via recursion [45] or

agentic workflows [185, 200]. Relatedly, changing the allocation of parameters to e.g. vocabulary versus non-vocabulary parameters is another avenue for improvement [170].

More data We train **OLMoE-1B-7B** for 5 trillion tokens, however, some recent dense models train significantly longer, such as Llama 3 with 15 trillion tokens [50]. To the best of our knowledge, there has been no large MoE that has been overtrained [57] as much as **OLMoE-1B-7B**. Specifically, taking the active parameters of **OLMoE-1B-7B**, our token multiplier [57] is around 5,000 (5T / 1B). There are likely benefits to training even longer, but to what degree overtraining is effective for MoEs and how it differs from dense models still requires more research [7].

Multimodal **OLMoE-1B-7B** is a text-only large language model, thus it cannot take inputs or produce outputs in other modalities like images or audio. This limits its utility for the large variety of multimodal use cases of such models [75, 165, 27, 82, 118, 134, 14, 47, 50]. There has been early work on open multimodal MoEs [124, 97, 94, 155, 111, 192] and we look forward to making future versions of **OLMoE** a part of that.

Multilingual We pretrain **OLMoE-1B-7B** on a predominantly English corpus and exclusively evaluate on English tasks. This may severely limit the usefulness of our model for research on non-English language models [107, 158, 221, 53, 163, 195]. While there has been work on training language-specific LMs [109, 55], it is more likely that as we add more data to build better future iterations of **OLMoE** we will mix in more non-English data due to data constraints [120]. This may make future **OLMoE** models perform better in non-English languages.