

# Domain Adaptation of Foundation LLMs for e-Commerce

Christian Herold\*   Michael Kozielski   Tala Bazazo   Pavel Petrushkov  
Patrycja Cieplicka   Dominika Basaj†   Yannick Versley   Seyyed Hadi Hashemi  
Shahram Khadivi

eBay Inc.

## Abstract

We present the e-Llama models: 8 billion and 70 billion parameter large language models that are adapted towards the e-commerce domain. These models are meant as foundation models with deep knowledge about e-commerce, that form a base for instruction- and fine-tuning. The e-Llama models are obtained by continuously pretraining the Llama 3.1 base models on 1 trillion tokens of domain-specific data.

We discuss our approach and motivate our choice of hyperparameters with a series of ablation studies. To quantify how well the models have been adapted to the e-commerce domain, we define and implement a set of multilingual, e-commerce specific evaluation tasks.

We show that, when carefully choosing the training setup, the Llama 3.1 models can be adapted towards the new domain without sacrificing significant performance on general domain tasks. We also explore the possibility of merging the adapted model and the base model for a better control of the performance trade-off between domains.

## 1 Introduction

Large Language Models (LLMs) have greatly improved the performance on most natural language tasks, and often show surprisingly good zero-shot generalization to new domains (Singhal et al., 2023). However, training on a specific target domain is often the means of choice to reach the best tradeoff in terms of scalability, domain knowledge, inference costs, and other factors.

While earlier approaches have trained domain-specific models from scratch (Beltagy et al., 2019; Alsentzer et al., 2019), the big effort of training competitive LLMs have meant that researchers and practitioners more commonly use continued pre-training (CPT), see e.g. Gururangan et al. (2020);

Ke et al. (2023), as compute requirements grew from using 1024 V100 GPUs over several days for RoBERTa as a larger BERT-like model (Liu et al., 2020) to the 16,000 H100 GPUs used for recent Llama-3 trainings (Dubey et al., 2024).

For e-commerce applications such as many seen at eBay, one could use existing pretrained models, such as Llama-3.1 (Dubey et al., 2024) for their use-cases. However, these models typically lack specific knowledge about the e-commerce domain.

Instead, we continue training the Llama base models on a large amount of e-commerce data. This way we introduce the domain specific knowledge into the model, while at the same time keeping the general capabilities of the model intact. This technique is known as ‘continued pretraining’ and the training setup has to be carefully balanced to prevent the model from degrading too much in performance on general domain tasks.

In Table 1, we compare recent continuous pre-training works in terms of the domain, the size of the models, as well as the amount of training data. As can be seen, our work is at a significantly larger scale than most existing works, either in terms of model size or in terms of tokens used for training, or both. We share our insights regarding large-scale model adaptation. In particular, we compare the model adaptation for models of different sizes and discuss the observed differences in behavior. We also explore the possibility of model merging to better control the trade-off between general- and domain-specific knowledge.

The rest of the paper is structured as follows. In Section 3 we discuss our data mixture and explain our methods of model evaluation with focus on e-commerce specific tasks. In Section 4 we explain our series of experiments to determine the optimal set of hyperparameters. In Section 5 we show the performance of the final models and discuss the possibility of model merging to better tune for different domains.

\*Correspondence author. Email: cherold@ebay.com.

†work done while at eBay

Study	Domain	Model Parameter Count	Total num Tokens
Minerva (Lewkowycz et al., 2022)	STEM	8B, 62B, 540B	26B-38.5B
MediTron (Chen et al., 2023)	Medicine	7B, 70B	46.7B
Code Llama (Rozière et al., 2023)	Code	7B, 13B, 34B, 70B	520B-1,000B
Llemma (Azerbayev et al., 2024)	Math	7B, 34B	50B-55B
DeepSeekMath (Shao et al., 2024)	Math	7B	500B
SaulLM-7B (Colombo et al., 2024b)	Law	7B	30B
SaulLM-54, 141B (Colombo et al., 2024a)	Law	54B, 141B	520B
HEAL (Yuan et al., 2024)	Medicine	13B	14.9B
Me-LLaMA (Xie et al., 2024)	Medicine	13B, 70B	129B
ClimateGPT (Thulke et al., 2024)	Climate	7B, 13B, 70B	4.2B
Nemotron (Parmar et al., 2024)	General	15B	1,000B
e-Llama (ours)	e-commerce	8B, 70B	1,000B

Table 1: Comparing the scale of recent continued pretraining works with our setting. Most existing works are at a significantly smaller scale, either in terms of model size or in terms of tokens used for training.

## 2 Related Work

The large cost of training LLMs from scratch has meant that continued pretraining is very attractive for adapting an existing LLM to new languages or domains. For example, Minixhofer et al. (2022) show that it is possible to reach competitive results for non-English languages by continuing the pre-training of RoBERTa and GPT-2 models. They start from an English model with tokenizer modification and reach scores on par with monolingual models trained from scratch on a multiple of the data used. In terms of adaptation to different domains, Gururangan et al. (2020) show that continued pretraining on a target domain helps a RoBERTa achieve better performance on tasks in that domain, even taking into account task-specific fine-tuning.

In terms of larger LLMs, Singhal et al. (2023) show that while zero-shot performance of PaLM finetuned on general-domain instruction data is surprisingly good on medical text, continued training on medical instruction data using parameter-efficient finetuning (PEFT) method can further improve these results. Lewkowycz et al. (2022) show that continued training on a mix of the original data and mathematical language from ArXiv and math web pages can boost PaLM’s performance on mathematical tasks. More recent papers address the problem of catastrophic forgetting in continued pretraining which can only partially be mitigated by using the original data in a portion of the continued pretraining mix: Ke et al. (2023) discuss masking updates to the neurons most instrumental for general-domain performance, and Wu et al.

(2024) show that growing the model by introducing additional layers, followed by only training these layers can avoid catastrophic forgetting. In contrast, newer work such as Ke et al. (2025) is more centered on an optimal data composition, proposing a mix of continued pretraining data and mixed-in instruction data.

In the e-commerce domain, Peng et al. (2024) as well as Li et al. (2024) focus exclusively on instruction tuning, while our work is the first to consider continued pretraining on domain-relevant data. For other domains, please refer to Table 1.

## 3 Setup, Data and Evaluation

### 3.1 Training Framework and Hardware

For training, we use the Megatron-LM framework from NVIDIA (Shoeybi et al., 2019; Narayanan et al., 2021). Training was conducted using 60 nodes, each having 8 NVIDIA H100 80GB GPUs (a total of 480 GPUs). The GPUs are connected via NVIDIA NVLink (intra-node) and InfiniBand (inter-node). The hardware is part of the eBay compute platform.

### 3.2 Data

Regarding training data, we mostly follow Herold et al. (2024). For general domain data, we use a mixture of web-crawled and smaller but more high quality datasets. We include 10% non-English general domain data in the data mix. Regarding the e-commerce domain, we employ several data sources. On the one hand, we utilize listings and product reviews from the eBay website, as has been

done by Herold et al. (2024). Furthermore, inspired by Lozhkov et al. (2024), we train an e-commerce classifier and use it to extract e-commerce specific examples from the Fineweb corpus (Penedo et al., 2024). We use this data for 20% of our e-commerce specific data mixture.

### 3.3 Evaluation

We perform evaluation both on general and e-commerce specific tasks. As a first benchmark, we calculate model perplexity on heldout datasets for general and e-commerce data.

#### General Domain

For evaluating the model capabilities on the general domain for the English language, we utilize the Natural Language Understanding (NLU) benchmark aggregates (in the following called **NLU En**) also used by Groeneveld et al. (2024) and Herold et al. (2024) and calculated using the EleutherAI LM Evaluation Harness (Gao et al., 2023). Furthermore, we utilize the ‘Open LLM Leaderboard 2’ (Fourrier et al., 2024) (in the following called **LLM Leaderboard En**) benchmark, which calculates a re-normalized average of the scores for the BBH (Suzgun et al., 2022), GPQA (Rein et al., 2023), MUSR (Sprague et al., 2024) and MMLU-PRO (Wang et al., 2024) benchmarks.<sup>1</sup> When we report model performance in Section 4, we average NLU En and LLM Leaderboard En scores. For the evaluation of the non-English, general domain NLU capabilities we use the same task aggregates as Herold et al. (2024) (in the following called **NLU non-En**). Furthermore, we utilize the ‘Open Multilingual LLM Leaderboard’ (Lai et al., 2023) (in the following called **LLM Leaderboard non-En**). In this work we focus on German, Spanish, French and Italian.

#### e-Commerce

Since existing work, like eCeLLM (Peng et al., 2024) and EcomGPT (Li et al., 2024) focuses on evaluation of instruction tuned models, we define a total of 5 novel e-commerce benchmarks for evaluation of foundation models. All tasks are strongly connected to relevant downstream tasks that we encounter in the e-commerce setting. They revolve around the listings on an e-commerce website, of which we consider title, category, price and a list

of aspect key-value pairs<sup>2</sup>. Below we list the tasks in detail:

1. **Aspect Prediction (AP)**: Given the title and category of a listing, as well as a specific aspect key, predict the corresponding aspect value.
2. **Aspect Prediction Multiple Choice (AP<sup>MC</sup>)**: Given 4 listings, of which 3 are corrupted by changing at least 1 aspect value, the model has to identify the correct listing.
3. **Price Prediction Multiple Choice (PP<sup>MC</sup>)**: Given 4 listings, of which 3 are corrupted by changing the price at which the item was sold, the model has to identify the listing with the correct selling price.
4. **Most Common Aspects (MCA)**: Given a category and an aspect key, the model has to predict the most common aspect values for that key.
5. **Most Common Aspects Multiple Choice (MCA<sup>MC</sup>)**: Given a category and an aspect key, the model is presented with 4 choices for the most common aspect value for that key and has to select the correct one.

We evaluate these tasks for English, German, Spanish, French and Italian. For all tasks, the final evaluation metric is accuracy. We give an example for each of the task in Appendix A.1.

In order to obtain a strong baseline, we perform a set of experiments where we optimize the number of few-shot examples for the base Llama-3 model, see Appendix A.2 for the details.

## 4 Finding the best Setup

In this section we discuss several series of experiments we performed to determine the best setup for continuously pretraining. For these studies we focus on the English language benchmarks, since we assume the non-English languages will follow the same trend. Since the 3.1 version of Llama was not released at the time, some experiments utilize Llama-3.0 models instead. The final models described in Section 5 are based on Llama-3.1.

<sup>1</sup>We exclude IFEval and MATH Lvl 5 benchmarks because the former is only useful for instruction-tuned models and the latter gives very low scores for the base models, especially for the 8B model variants.

<sup>2</sup>An example for an aspect key could be ‘Brand’ and a possible aspect-value in this case could be ‘Nike’.

$LR_{max}$	ppl ( $\downarrow$ )		benchmark ( $\uparrow$ )	
	e-com.	general	e-com.	general
Llama-3.0	7.28	8.38	45.9	44.1
3.0e-5	2.03	6.43	59.9	42.0
3.0e-4	2.01	6.48	58.6	40.5
3.0e-3	2.15	7.70	50.6	34.5

Table 2: Effect of the maximum learning rate of the continued pretraining (1 trillion tokens) of Llama-3.0 8B on the final model performance. Llama-3.0 used  $LR_{max}=3.0e-4$ .

#### 4.1 Learning Rate

Maybe the most important hyperparameter to consider is the maximum learning rate  $LR_{max}$  of the continued pretraining. Meta have used a  $LR_{max}$  of  $3.0e-4$  and  $1.5e-4$  for their training of Llama-3.1 8B and 70B respectively (Dubey et al., 2024). However, using the same maximum learning rate for continued pretraining might not yield the best results as the model might forget too much information from the previous training or, on the contrary, the model might not learn enough from the new data mixture (Gupta et al., 2023; Ibrahim et al., 2024).

There are mainly 2 paradigms in existing work: (i) use the same  $LR_{max}$  as for the original pretraining (Rozière et al., 2023; Chen et al., 2023; Shao et al., 2024), or (ii) use a smaller value, typically around 10% of the original  $LR_{max}$  (Azerbayev et al., 2024; Lewkowycz et al., 2022; Colombo et al., 2024a; Yuan et al., 2024; Thulke et al., 2024; Xie et al., 2024; Parmar et al., 2024).

We perform a set of experiments to determine the best maximum learning rate. Since the impact of the learning rate might significantly depend on the amount of data used in training, we decide to compare training runs utilizing the full 1 trillion tokens of data (50% e-commerce ratio). In all cases, the learning rate decays over the course of the full training with a cosine scheduling to the minimum learning rate of  $3.0e-6$ . We compare the final model performance in terms of perplexity on the heldout test sets, as well as general (average of NLU En and LLM Leaderboard) and e-commerce benchmarks. The results can be found in Table 2.

In terms of perplexity, we find that a higher learning rate leads to a slightly better score on the new domain. However, these improvements do not translate to a better score on the e-commerce spe-

% e-com	ppl ( $\downarrow$ )		benchmark ( $\uparrow$ )	
	e-com.	general	e-com.	general
Llama-3.0	7.28	8.38	45.9	44.1
10	2.75	6.87	55.6	43.2
25	2.59	6.92	56.7	43.1
50	2.47	7.00	57.5	43.3
75	2.40	7.15	57.6	43.2

Table 3: Effect of the amount of e-commerce data in the continued pretraining (30 billion tokens) of Llama-3.0 8B on the final model performance.




Figure 1: Perplexity as a function of the input sequence length for the 8B (e)-Llama-3.0/3.1 models. The 3.0 variants can not handle context sizes much longer than 8k, since they have never seen these lengths in training.

cific benchmarks. At the same time, a higher learning rate leads to more degradation on the general domain benchmarks. This might be an indication that our general domain data mix is maybe a bit lower quality than what has been used by Meta in the Llama-3 pretraining. In the end, we decide to use an  $LR_{max}$  that is 10% of the maximum learning rate used in pretraining, i.e. **3.0e-5 for the 8B model and 1.5e-5 for the 70B model**.

#### 4.2 Data Weighting

While we want the model to learn about the new domain, at the same time we want to avoid the effect of catastrophic forgetting. To combat this, it is common to include some percentage of general domain examples in the data mixture (sometimes called ‘replay examples’). Most existing works use only up to 15% of general data in their mixture (Azerbayev et al., 2024; Lewkowycz et al., 2022; Rozière et al., 2023; Chen et al., 2023; Colombo et al., 2024b,a) with the exception of Yuan et al. (2024) who use 35%. However, we have reason to






Figure 2: Evolution of the model performance over the course of the training. Left: perplexity on heldout data sets in the general/e-commerce domain. Center/Right: Evolution of downstream model performance on English tasks for 8B/70B model.

believe that in our specific case, a higher ratio of general domain data might be advisable, because our in-domain e-commerce data comes from a very different distribution.

Following Ibrahim et al. (2024) we perform experiments on a limited amount of training tokens (ca 30 billion) with a varying ratio of e-commerce data ( $LR_{max}=3.0e-5$ ). The results can be found in Table 3.

As expected, we see that with a higher percentage of e-commerce data, the perplexity and downstream performance for the e-commerce related tasks is improving, although with diminishing returns when going above 50%. At the same time, perplexity on the general domain data is getting worse, but this does not effect the model scores on the general domain benchmarks. As mentioned before, the reason for this is most likely the different distributions of the general domain data we are using vs the one that was used in pretraining. In the end, we decide to continue pretraining with an **e-commerce percentage of 50% in our data mix**.

### 4.3 Context Size

Finally, we explore the effect of the continued pretraining on the context size of the model. While Llama-3.0 has a context size of 8k, the Llama-3.1 models have a much larger context size of 128k. Ideally we would like to continue the pretraining with the same large context size, but this introduces several challenges. First, the vast majority of our

training examples both for general and e-commerce domain are shorter than 1k tokens. Additionally, increasing the context size makes it harder to train the model efficiently due to the quadratic computational complexity of the transformer model. There exist methods to mitigate the latter issue, like the context parallel training approach (Fang and Zhao, 2024) but when applying said approach, we find that this still introduces too much computational overhead and significantly slows down the training. We decide to continuously train both Llama-3.0 8B and 3.1 8B with 8k context size and study the effects this has on the models. In Figure 1 we calculate the perplexity of the 8B base and e-Llama models as a function of the input sequence length for a general domain heldout test set.

All models exhibit nearly identical performance for inputs smaller than 8k. Unsurprisingly, the 3.0 variants can not handle inputs larger than 8k at all. The e-Llama model that is based on Llama-3.1 exhibits a much better understanding of longer sequences, even though it has not seen any sequences longer than 8k in the continued pretraining. We can conclude that the model retains most of its ability to handle longer sequences. We do see some degradation for even longer input lengths, but this is an acceptable trade-off for us. We therefore decide to **perform the continued pretraining with a context size of 8k**.

Model	general domain benchmarks (↑)				e-commerce benchmarks (↑)					
	En		non-En		AP	AP <sup>MC</sup>	En PP <sup>MC</sup>	MCA	MCA <sup>MC</sup>	non-En avg.
	NLU	Lead.	NLU	Lead.						
<b>8B</b>										
Llama-3.1	71.8	17.2	54.1	43.2	36.5	61.8	50.1	27.4	55.3	35.8
e-Llama	71.6	12.6	54.0	42.4	54.9	74.9	59.6	37.8	67.4	46.8
<b>70B</b>										
Llama-3.1	76.6	29.7	58.5	55.2	42.8	66.3	59.3	35.2	61.9	40.4
e-Llama	76.3	28.7	59.2	55.4	59.2	79.5	65.7	49.9	71.5	52.8

Table 4: Final performance of the e-Llama 8B/70B models on general domain and e-commerce specific evaluation benchmarks.

## 5 e-Llama

In this section we discuss the training and performance of the final e-Llama 8B and 70B models.

### 5.1 Training

Our setup mostly follows [Herold et al. \(2024\)](#) while taking into account our findings from Section 4. In particular we use cosine Learning Rate (LR) scheduling with warmup, a batch-size of ca. 11.8 million tokens and 85k total update steps.

In Figure 2, we show the evolution of the 8B/70B model performance over the course of the training. We see that the perplexity on the general domain data is decreasing for both 8B and 70B model. At the same time, the gap between 8B and 70B stays constant throughout the training. This indicates that while the distribution of our general domain data is different from the original one, the complexity of the data might be similar. Perplexity on the e-commerce data is also decreasing but at a much faster rate. In the end, the difference in terms of e-commerce perplexity for 8B and 70B model is much smaller than for the base models.

In terms of downstream performance, we find that the 8B model seems to quickly become saturated and performance is no longer increasing after 20% of the training. The 70B model on the other hand recovers much better on the general domain tasks, while also continuously improving in the e-commerce domain. We think this might be due to the much larger model size, that allows the model to better incorporate new information without catastrophic forgetting. Also, the smaller learning rate for the 70B model might have played a role here.

### 5.2 Final models

In Table 4 we show the final model performance of e-Llama 8B/70B in comparison to the Llama-3.1 base models.

On the general domain NLU benchmarks, the e-Llama models perform the same as the base Llama-3.1 models. On the more challenging LLM Leaderboard tasks, we see some performance degradation, especially for the smaller 8B model variant. We think this might be due to a combination of smaller model size and a different data distribution of our general domain data compared to what has been used at Meta.

On the e-commerce benchmarks, the e-Llama models improve relative to the Llama-3.1 base models by around 25% on English and by around 30% on non-English benchmarks on average.<sup>3</sup> Interestingly, the gap between 8B and 70B variant for Llama base model and for e-Llama is roughly the same for the e-commerce tasks, even though in terms of perplexity the e-Llama models are closer together (compare left side of Figure 2). This once again highlights that perplexity and downstream performance do not always follow the same trend.

### 5.3 Model Merging

The last topic we want to discuss is how to better align the trade-off between general and domain-specific performance. Lets assume we want to have the best performing model on the e-commerce domain, but we can not allow the general domain performance to drop below a certain threshold on the general domain tasks. As can be seen from the experiments in Figure 2 and Table 3, reducing the percentage or amount of e-commerce train-

<sup>3</sup>The individual language scores for non-English can be found in Table 5 in the Appendix.




Figure 3: Model merging: 8B Model performance on English general and e-commerce benchmarks as a function of the weight of the e-Llama model parameters vs the base Llama-3.1 model parameters.

ing data does not allow to make very precise forecasts of final model performance. Instead, we utilize a technique called ‘model merging’ (Wortsman et al., 2022), where we simply average all parameters of the base Llama-3.1 model checkpoint and our final e-Llama model checkpoint. In Figure 3 we show how the performance of the resulting model changes as a function of the individual model weights.

The performance for both general and e-commerce domain follow an almost linear trend. This allows for a very precise tuning of the final model performance and has the additional advantage that the model merging is not compute intensive at all.

## 6 Conclusion

We have discussed our efforts to adapt the Llama-3.1 8B and 70B parameter base models towards the e-commerce domain. In order to evaluate the model capabilities in the e-commerce setting, we design and implement a set of multilingual, e-commerce specific evaluation benchmarks. Through a series of experiments, we determine the best experimental setting for our use-case. We show that the models can be adapted well towards the new domain with limited degradation on general domain performance. Furthermore, we highlight that with model merging, we can very precisely tune the final model performance.

## 7 Limitations

The present work has several limitations: (i) We focus on a single domain only, namely e-commerce. (ii) We focus on non-instruction-tuned foundation

models only. A logical improvement is to consider instruction tuning as an additional part in the pipeline. (iii) While we try to define a comprehensive set of evaluations for the e-commerce domain, the diversity and quantity of evaluations could be further improved. (iv) Finally, in this work we focus solely on the Llama family of models. Future work should explore further open source models.

## References

- Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical BERT embeddings](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen Marcus McAleer, Albert Q. Jiang, Jia Deng, Stella Biderman, and Sean Welleck. 2024. [Llemma: An open language model for mathematics](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3615–3620, Hong Kong, China. Association for Computational Linguistics.
- Zeming Chen, Alejandro Hernández-Cano, Angelika Romanou, Antoine Bonnet, Kyle Matoba, Francesco Salvi, Matteo Pagliardini, Simin Fan, Andreas Köpf, Amirkeivan Mohtashami, Alexandre Sallinen, Alireza Sakhaeirad, Vinitra Swamy, Igor Krawczuk, Deniz Bayazit, Axel Marmet, Syrielle Montariol, Mary-Anne Hartley, Martin Jaggi, and Antoine Bosselut. 2023. [MEDITRON-70B: scaling medical pretraining for large language models](#). *CoRR*, abs/2311.16079.
- Pierre Colombo, Telmo Pires, Malik Boudiaf, Rui Melo, Dominic Culver, Sofia Morgado, Etienne Malaboeuf, Gabriel Hautreux, Johanne Charpentier, and Michael Desa. 2024a. [Saullm-54b & saullm-141b: Scaling up domain adaptation for the legal domain](#). *CoRR*, abs/2407.19584.
- Pierre Colombo, Telmo Pessoa Pires, Malik Boudiaf, Dominic Culver, Rui Melo, Caio Corro, André F. T. Martins, Fabrizio Esposito, Vera Lúcia Raposo, Sofia Morgado, and Michael Desa. 2024b. [Saullm-7b: A pioneering large language model for law](#). *CoRR*, abs/2403.03883.
- 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, and 82 others. 2024. *The llama 3 herd of models*. *CoRR*, abs/2407.21783.
- Jiarui Fang and Shangchun Zhao. 2024. *USP: A unified sequence parallelism approach for long context generative AI*. *CoRR*, abs/2405.07719.
- Clémentine Fourier, Nathan Habib, Alina Lozovskaya, Konrad Szafer, and Thomas Wolf. 2024. Open llm leaderboard v2. [https://huggingface.co/spaces/open-llm-leaderboard/open\\_llm\\_leaderboard](https://huggingface.co/spaces/open-llm-leaderboard/open_llm_leaderboard).
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, and 5 others. 2023. *A framework for few-shot language model evaluation*.
- 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, and 24 others. 2024. *Olmo: Accelerating the science of language models*. *CoRR*, abs/2402.00838.
- Kshitij Gupta, Benjamin Thérien, Adam Ibrahim, Mats L. Richter, Quentin Anthony, Eugene Belilovsky, Irina Rish, and Timothée Lesort. 2023. *Continual pre-training of large language models: How to (re)warm your model?* *CoRR*, abs/2308.04014.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. *Don’t stop pretraining: Adapt language models to domains and tasks*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Christian Herold, Michael Kozielski, Leonid Eki-mov, Pavel Petrushkov, Pierre-Yves Vandenbussche, and Shahram Khadivi. 2024. *Lilium: ebay’s large language models for e-commerce*. *CoRR*, abs/2406.12023.
- Adam Ibrahim, Benjamin Thérien, Kshitij Gupta, Mats L. Richter, Quentin Gregory Anthony, Eugene Belilovsky, Timothée Lesort, and Irina Rish. 2024. *Simple and scalable strategies to continually pre-train large language models*. *Trans. Mach. Learn. Res.*, 2024.
- Zixuan Ke, Zixuan Ke1, Yijia Shao, Haowei Lin, Tatsuya Konishi, Gyuhak Kim, and Bing Liu. 2023. *Continual pre-training of language models*. In *The Eleventh International Conference on Learning Representations (ICLR 2023)*.
- Zixuan Ke, Yifei Ming, Xuan-Phi Nguyen, Caiming Xiong, and Shafiq Joty. 2025. *Demystifying domain-adaptive post-training for financial llms*. *Preprint*, arXiv:2501.04961.
- Viet Lai, Nghia Trung Ngo, Amir Pouran Ben Veyseh, Franck Dernoncourt, and Thien Huu Nguyen. 2023. *Open multilingual llm evaluation leaderboard*.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay V. Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. 2022. *Solving quantitative reasoning problems with language models*. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Yangning Li, Shirong Ma, Xiaobin Wang, Shen Huang, Chengyue Jiang, Haitao Zheng, Pengjun Xie, Fei Huang, and Yong Jiang. 2024. *Ecomgpt: Instruction-tuning large language models with chain-of-task tasks for e-commerce*. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, pages 18582–18590. AAAI Press.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2020. *Roberta: A robustly optimized bert pretraining approach*. In *The Eighth International Conference on Learning Representations (ICLR 2020)*.
- Anton Lozhkov, Loubna Ben Allal, Leandro von Werra, and Thomas Wolf. 2024. *Fineweb-edu*.
- Benjamin Minixhofer, Fabian Paischer, and Navid Rekasaz. 2022. *WECHSEL: Effective initialization of subword embeddings for cross-lingual transfer of monolingual language models*. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3992–4006, Seattle, United States. Association for Computational Linguistics.
- Deepak Narayanan, Mohammad Shoeybi, Jared Casper, Patrick LeGresley, Mostofa Patwary, Vijay Korthikanti, Dmitri Vainbrand, Prethvi Kashinkunti, Julie Bernauer, Bryan Catanzaro, Amar Phanishayee, and Matei Zaharia. 2021. *Efficient large-scale language model training on GPU clusters using megatron-lm*. In *International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2021, St. Louis, Missouri, USA, November 14-19, 2021*, page 58. ACM.



- Jupinder Parmar, Sanjeev Satheesh, Mostofa Patwary, Mohammad Shoeybi, and Bryan Catanzaro. 2024. [Reuse, don't retrain: A recipe for continued pretraining of language models](#). *CoRR*, abs/2407.07263.
- Guilherme Penedo, Hynek Kydlíček, Loubna Ben al-lal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. 2024. [The fineweb datasets: Decanting the web for the finest text data at scale](#). *Preprint*, arXiv:2406.17557.
- Bo Peng, Xinyi Ling, Ziru Chen, Huan Sun, and Xia Ning. 2024. [ecellm: Generalizing large language models for e-commerce from large-scale, high-quality instruction data](#). *CoRR*, abs/2402.08831.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Driani, Julian Michael, and Samuel R. Bowman. 2023. [Gpqa: A graduate-level google-proof q&a benchmark](#). *Preprint*, arXiv:2311.12022.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton-Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, and 6 others. 2023. [Code llama: Open foundation models for code](#). *CoRR*, abs/2308.12950.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *CoRR*, abs/2402.03300.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. [Megatron-lm: Training multi-billion parameter language models using model parallelism](#). *CoRR*, abs/1909.08053.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Abubakr Babiker, Nathanael Schärli, Aakanksha Chowdhery, Philip Mansfield, Dina Demner-Fushman, and 13 others. 2023. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180.
- Zayne Sprague, Xi Ye, Kaj Bostrom, Swarat Chaudhuri, and Greg Durrett. 2024. [Musr: Testing the limits of chain-of-thought with multistep soft reasoning](#). *Preprint*, arXiv:2310.16049.
- 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](#). *Preprint*, arXiv:2210.09261.
- David Thulke, Yingbo Gao, Petrus Pelsers, Rein Brune, Richa Jalota, Floris Fok, Michael Ramos, Ian van Wyk, Abdallah Nasir, Hayden Goldstein, Taylor Tragemann, Katie Nguyen, Ariana Fowler, Andrew Stanco, Jon Gabriel, Jordan Taylor, Dean Moro, Evgenii Tsymbalov, Juliette de Waal, and 7 others. 2024. [Climategpt: Towards AI synthesizing interdisciplinary research on climate change](#). *CoRR*, abs/2401.09646.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhui Chen. 2024. [Mmlu-pro: A more robust and challenging multi-task language understanding benchmark](#). *Preprint*, arXiv:2406.01574.
- Mitchell Wortsman, Gabriel Ilharco, Samir Yitzhak Gadre, Rebecca Roelofs, Raphael Gontijo Lopes, Ari S. Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. 2022. [Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 23965–23998. PMLR.
- Chengyue Wu, Yukang Gan, Yixiao Ge, Zeyu Lu, Jiahao Wang, Ye Feng, Ying Shan, and Ping Luo. 2024. [LLaMA pro: Progressive LLaMA with block expansion](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6518–6537, Bangkok, Thailand. Association for Computational Linguistics.
- Qianqian Xie, Qingyu Chen, Aokun Chen, Cheng Peng, Yan Hu, Fongci Lin, Xueqing Peng, Jimin Huang, Jeffrey Zhang, Vipina Kuttichi Keloth, Xingyu Zhou, Huan He, Lucila Ohno-Machido, Yonghui Wu, Hua Xu, and Jiang Bian. 2024. [Me llama: Foundation large language models for medical applications](#). *CoRR*, abs/2402.12749.
- Dong Yuan, Eti Rastogi, Gautam Naik, Sree Prasanna Rajagopal, Sagar Goyal, Fen Zhao, Bharath Chintagunta, and Jeff Ward. 2024. [A continued pretrained LLM approach for automatic medical note generation](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Short Papers, NAACL 2024, Mexico City, Mexico, June 16-21, 2024*, pages 565–571. Association for Computational Linguistics.

## A Appendix

### A.1 Examples for e-commerce tasks

Here, we give an example for each of the e-commerce tasks described in Section 3.3. All examples are for the English language. For languages other than English, the prompt stays the same, but the item-specific attributes like title are in the corresponding language.

#### AP

The model has to predict the most probable continuation of the following text input:

For an e-commerce website, under the category "Video Games & Consoles:Video Games", the listing with the title "Dark Souls III (Sony PlayStation 4)" has the following aspect key-value pairs:

Rating:

#### AP<sup>MC</sup>

The model is used to independently score the following 4 text sequences, has to give the highest probability to the correct sequence (first one).

For an e-commerce website, the listing with the title "Dark Souls III (Sony PlayStation 4)" has the following aspect key-value pairs associated with it:

Rating: M - Mature

For an e-commerce website, the listing with the title "Dark Souls III (Sony PlayStation 4)" has the following aspect key-value pairs associated with it:

Rating: E - Everyone

For an e-commerce website, the listing with the title "Dark Souls III (Sony PlayStation 4)" has the following aspect key-value pairs associated with it:

Rating: T - Teen

For an e-commerce website, the listing with the title "Dark Souls III (Sony PlayStation 4)" has the following aspect key-value pairs associated with it:

Rating: AO - Adults Only

#### PP<sup>MC</sup>

The model is used to independently score the following 4 text sequences, has to give the highest probability to the correct sequence (first one).

For the listing with the title "Authentic Louis Vuitton Monogram Empreinte Bastille PM 2Way Tote Bag Black 9281E", the final selling price was \$816.00.

For the listing with the title "Authentic Louis Vuitton Monogram Empreinte Bastille PM 2Way Tote Bag Black 9281E", the final selling price was \$81.60.

For the listing with the title "Authentic Louis Vuitton Monogram Empreinte Bastille PM 2Way Tote Bag Black 9281E", the final selling price was \$204.00.

For the listing with the title "Authentic Louis Vuitton Monogram Empreinte Bastille PM 2Way Tote Bag Black 9281E", the final selling price was \$1632.00.

#### MCA

The model has to predict the most probable continuation of the following text input:

For an e-commerce website, under the category "Clothing, Shoes & Accessories:Women:Women's Clothing:Coats, Jackets & Vests", the following are the most common aspect values for the aspect key "Outer Shell Material":

#### MCA<sup>MC</sup>

The model is used to independently score the following 4 text sequences, has to give the highest probability to the correct sequence (first one).

For an e-commerce website, under the category "Cameras & Photo:Digital Cameras", the most common aspect value for the aspect key "Brand" is "Canon".

For an e-commerce website, under the category "Cameras & Photo:Digital Cameras", the most common aspect value for the aspect key "Brand" is "Fujifilm".

For an e-commerce website, under the category "Cameras & Photo:Digital




Figure 4: Llama-3 8B model performance on the 5 e-commerce evaluation tasks as a function of the number of few-shot examples provided in the prompt.

Cameras", the most common aspect value for the aspect key "Brand" is "PENTAX".

For an e-commerce website, under the category "Cameras & Photo:Digital Cameras", the most common aspect value for the aspect key "Brand" is "Nikon".

## A.2 Optimizing the Few-Shot Setup

For the base Llama-3 8B model, we prompt the model for each of the above tasks with up to 20 few-shot examples. The results can be seen in Figure 4.

We find that for the PP<sup>MC</sup> task, few-shot prompting does not significantly improve the model performance, therefore in the following we use **0-shot** evaluation for this task. Since we have already a quite high score for the AP<sup>MC</sup> task, and there is only limited information to be gained from the few-shot examples, we decide to use **1-shot** evaluation for this task. For both AP and MCA<sup>MC</sup> we see improvements with more few-shot examples. Therefore we end up using **5-shot** evaluation for these tasks. Finally for MCA we have quite low scores overall, and the model seems to benefit from more few-shot examples. Therefore we end up using **20-shot** evaluation for this task.

## A.3 Non-English Benchmark Scores

Model	e-commerce benchmarks ( $\uparrow$ )			
	De	Fr	It	Es
<b>8B</b>				
Llama-3.1	35.4	35.0	35.0	37.6
e-Llama	47.5	47.0	46.0	46.7
<b>70B</b>				
Llama-3.1	39.2	40.6	40.7	41.2
e-Llama	52.7	52.2	54.5	52.0

Table 5: Final performance of the e-Llama 8B/70B models on language-specific, e-commerce evaluation benchmarks.