

ENHANCING MULTILINGUAL LLM PRETRAINING WITH MODEL-BASED DATA SELECTION

Bettina Messmer*, Vinko Sabolčec*, Martin Jaggi

School of Computer and Communication Sciences

MLO, EPFL

Lausanne, Switzerland

{bettina.messmer, vinko.sabolcec, martin.jaggi}@epfl.ch

ABSTRACT

Dataset curation has become a basis for strong large language model (LLM) performance. While various rule-based filtering heuristics exist for English and multilingual datasets, model-based filtering techniques have primarily focused on English. To address the disparity stemming from limited research on non-English languages, we propose a model-based filtering framework for multilingual datasets that aims to identify a diverse set of structured and knowledge-rich samples. Our approach emphasizes transparency, simplicity, and efficiency, leveraging Transformer- and FastText-based classifiers to ensure the broad accessibility of our technique and data. We conduct comprehensive ablation studies on the FineWeb-2 web crawl dataset across diverse language families, scripts, and resource availability to demonstrate the effectiveness of our method. Using a 1B-parameter Llama model trained on 70B and 119B tokens, our approach can match the baseline MMLU score with as little as 15% of the training tokens, while also improving across other benchmarks. These findings provide strong evidence for the generalizability of our approach to other languages. As a result, we extend our framework to 20 languages for which we will release the refined pretraining datasets.

1 INTRODUCTION

Large Language Models (LLMs) have demonstrated impressive performance improvements when trained on increasingly larger datasets and model sizes (Brown et al., 2020). While Brown et al. (2020) already observed the importance of using a cleaned version of Common Crawl for improved performance, the high cost of LLM training has further motivated research into better pretraining quality filters.

Deduplication and heuristic-based dataset cleaning have become standard practices in data curation (Rae et al., 2021; Raffel et al., 2020; De Gibert et al., 2024). These quality filters are often complemented by additional filters, such as the removal of personally identifiable information (PII) (Penedo et al., 2024a) or model-based toxicity filtering (Soldaini et al., 2024). Recently, model-based filtering has also emerged as a promising method for quality filtering. The release of FineWeb-Edu (Penedo et al., 2024a) demonstrated that pretraining on just 10% of the tokens (38B) from an English dataset filtered using a model-based approach can achieve performance comparable to models trained on 350B tokens of unfiltered data. Moreover, when trained on equivalent amounts of data, this model largely outperforms the baseline. Concurrently, the release of DCLM (Li et al., 2024b) showed that competitive performance can be achieved using a simple and efficient model-based approach, namely a FastText (Joulin et al., 2017) classifier trained on a carefully selected training dataset.

However, these recent advances have primarily focused on English data. This emphasis risks further widening the disparity in LLM performance between languages, as less than half of internet content is written in English¹. To address this concern, we aim to extend model-based filtering frameworks to multilingual datasets. While model perplexity-based filtering is commonly applied to multilingual

*Equal contribution

¹w3techs.com/technologies/overview/content_language

datasets (Wenzek et al., 2019; Laurençon et al., 2022; Nguyen et al., 2023), the current state-of-the-art, FineWeb-2 (Penedo et al., 2024c), primarily relies on heuristic-based filters. In this work, we focus on model-based filtering with a quality definition that emphasizes: 1) structured data and 2) knowledge-rich data samples, to enhance multilingual pretraining datasets.

To achieve this, we leverage embedding-based classification models. Firstly, we adopt the FastText quality filtering approach from DCLM to develop a unified framework for multilingual datasets that span diverse language families, scripts, and resource availability, focusing on Chinese, German, French, Arabic, and Danish as representative languages for our experiments. Additionally, we extend this embedding-based approach by incorporating Transformer (Vaswani et al., 2023) embeddings, specifically XLM-RoBERTa (Conneau et al., 2020), for filtering.

In summary, our contributions are as follows:

- We propose a transparent, simple, and unified framework for multilingual model-based filtering at web scale, enabling data curation across diverse language families, scripts and resource availability.
- We present comprehensive per-language ablation studies of embedding-based multilingual quality filtering on top of the FineWeb-2 dataset (Penedo et al., 2024c), achieving performance comparable to the baseline while using as little as 15% of the tokens. We additionally analyze the impact of dataset contamination and multilingual LLM training.
- We evaluate the impact of different training datasets for data selection classifiers on the downstream performance of LLMs.
- We release the codebase² and the refined pretraining dataset covering 20 languages³, filtered using our proposed framework, to advance multilingual language modeling.

2 RELATED WORK

Data Curation. In order to pretrain LLMs on a large amount of diverse texts, Common Crawl⁴ is often used as the base dataset. However, early works already observed that performing quality filtering on Common Crawl is crucial for model performance (Brown et al., 2020). There exist various data curation approaches, such as deduplication (Lee et al., 2022), PII removal (Subramani et al., 2023), or toxicity filtering (Arnett et al., 2024). Another important aspect is quality filtering of the documents. For this, the definition of quality is an important aspect. A common approach is to use heuristics to remove documents outside of the target distribution, such as filtering based on average word length, existence of punctuation, or document length (Rae et al., 2021; Raffel et al., 2020). Another approach is to define model-based filters, where research has focused on perplexity measure of the text (Wenzek et al., 2019; Marion et al., 2023; Ankner et al., 2024), distributional similarity measures (Brown et al., 2020; Xie et al., 2023; Li et al., 2024b) and LLM-based quality assessment (Gunasekar et al., 2023; Wettig et al., 2024; Sachdeva et al., 2024; Penedo et al., 2024a).

In this work, we build upon previous curated datasets based on heuristic filtering, specifically FineWeb-2 (Penedo et al., 2024c), and focus on model-based distributional similarity filtering for structured and knowledge-rich documents relying on textual embedding representation.

Curated English datasets. One of the early curated datasets was C4 (Raffel et al., 2020), followed by MassiveText (Rae et al., 2021). RefinedWeb (Penedo et al., 2023) was an important step forward, demonstrating that filtered web data can outperform selected high-quality data sources. While these datasets have not been made fully publicly available, their filtering techniques have been expanded upon in recent fully public datasets, such as Dolma (Soldaini et al., 2024), FineWeb, and FineWeb-Edu (Penedo et al., 2024a). While FineWeb primarily relies on filter heuristics for data quality, Dolma adopts model perplexity filtering. FineWeb-Edu takes model-based filtering a step further and relies on LLM-based quality assessment. Similarly, a concurrent work, DCLM, has achieved competitive performance using FastText (Joulin et al., 2017) classifier trained on a carefully selected training dataset. In this work we adapt and extend this approach to the multilingual context.

²github.com/epfml/fineweb2-hq

³huggingface.co/datasets/epfml/FineWeb2-HQ

⁴commoncrawl.org

Curated Multilingual Datasets. Analogously to the English datasets, there have been efforts in the multilingual space. An influential work has been CCNet (Wenzek et al., 2019), whose language identification and model perplexity filter for data quality has been re-used in later datasets. Again, while CCNet was not published directly, but rather provided the tools for data cleaning, RedPajama (Together Computer, 2023) is a prominent multilingual dataset relying on these filtering techniques. While RedPajama offers data in 5 European languages, other datasets, such as OSCAR (Ortiz Suárez et al., 2019; Abadji et al., 2021; Abadji et al., 2022), mC4 (Xue et al., 2021), ROOTS (Laurençon et al., 2022), MADLAD-400 (Kudugunta et al., 2023), CulturaX (Nguyen et al., 2023), and HPLT (de Gibert et al., 2024), focus on expanding beyond, spanning a variety of language families and scripts. While they offer refined datasets for hundreds of languages, FineWeb-2 (Penedo et al., 2024c) pushes the limit to thousands of languages and further improves the performance. Our work also focuses on filtering quality samples across various language families and scripts. However, we limit our scope to 20 languages, as the number of documents drops quickly and there is trade-off between retaining a sufficient number of pretraining tokens and ensuring data quality (Muennighoff et al., 2023; Held et al., 2025). In our results, we observe the greatest benefits using stricter data filtering.

Multilingual Embedding Models. Early word embedding models like Word2Vec (Mikolov et al., 2013) and GloVe (Pennington et al., 2014) lacked contextual understanding. FastText (Bojanowski et al., 2017) built upon them and improved performance by incorporating subword information. Transformer (Vaswani et al., 2023) models like BERT (Devlin et al., 2019) and GPT (Radford et al., 2018) then revolutionized the field with context-aware embeddings. Multilingual models like mBERT, XLM (Lample & Conneau, 2019), and XLM-RoBERTa (Conneau et al., 2020) further advanced cross-lingual understanding, with recent open-source LLMs pushing performance even higher (Llama Team, 2024; Mistral AI, 2025). Using such models, documents as well as representative samples can be mapped into a shared embedding space to estimate their similarity. Focusing on transparency, simplicity and efficiency in our work, we use FastText and XLM-RoBERTa for our filtering, and analyze the trade-off between computational complexity and filtering performance.

Multilingual Evaluation. Evaluating LLMs requires diverse benchmarks testing linguistic and cognitive abilities like reading comprehension, reasoning, and knowledge. While English benchmarks like MMLU (Hendrycks et al., 2020) and ARC (Clark et al., 2018) exist, other languages often use translations from English, e.g., XNLI (Conneau et al., 2018) and machine-translated version of MMLU (Lai et al., 2023). However, translations can be problematic, failing to capture cultural nuances or introducing "translationese" (Romanou et al., 2024). Recent work by Romanou et al. (2024); Singh et al. (2024a) emphasizes the need for culturally sensitive, natively collected benchmarks. Task difficulty and task formulation also impact model performance when trained for shorter durations (Kydliček et al., 2024). In our work, we follow the recent evaluation tasks selection and methodology by Kydliček et al. (2024) to assess our model-based filtering approaches across multiple languages.

3 METHODS

In this work, we present our model-based filtering approaches. Our methodology is structured into two key components: 1) we select suitable training datasets, aiming to identifying a diverse set of structured and knowledge-rich samples and 2) we describe the different models, namely FastText and Transformer embedding-based filters, used to capture and leverage these characteristics.

3.1 CLASSIFIER TRAINING DATASET

Representative Sample Selection. Our goal is to identify a diverse set of structured and knowledge-rich samples, especially within a multilingual context. We define two criteria for our training datasets: 1) the samples must be informative and well-structured and 2) the datasets must be available in multiple languages. While some multilingual benchmark datasets meet these criteria precisely, it is important to note that we do not train the LLM directly on this data. Instead, we train a proxy model to assess pretraining data quality. Nevertheless, we must remain cautious about potentially increased pretraining data contamination stemming from this approach, as discussed in Section 4.2.6.

Based on our criteria, we selected the following datasets as representative examples.

- **Aya Collection.** A prompt completion dataset comprising $\sim 514\text{M}$ samples covering a wide variety of tasks, generated using instruction-style templates in 101 languages (Singh et al., 2024b).
- **Aya Dataset.** Human-annotated instruction fine-tuning dataset consisting of $\sim 202\text{K}$ prompt-completion pairs in 65 languages (Singh et al., 2024b).
- **MMLU.** Originally for English language, the dataset contains $\sim 14\text{K}$ multiple-choice knowledge questions in diverse subjects and areas (Hendrycks et al., 2020). Multilingual version was translated into 14 languages by professional translators (OpenAI, 2024).
- **OpenAssistant-2.** The dataset contains $\sim 14\text{K}$ user-assistant conversations with multiple messages in 28 languages (Fischer et al., 2024).
- **Include-Base-44.** Multiple-choice questions focused on general and regional knowledge, as well as reasoning, extracted from academic and professional exams. Spanning 44 languages, it includes a total of $\sim 23\text{K}$ samples (Romanou et al., 2024).

Representative Sample Collection. *MMLU* and *Include-Base-44* are highly curated benchmark datasets, containing structured, knowledge-rich samples. The *Aya Dataset* is human-curated, while *OpenAssistant-2* is partially human-curated and partially generated by large language models (LLMs). In contrast, the *Aya Collection* consists of various AI-generated samples without quality guarantee, though it represents the largest and most multilingual of the five.

To address this quality difference, we create two *Multilingual Knowledge Collection (MKC)* configurations:

- **MKC:** Includes *Include-Base-44*, *OpenAssistant-2*, *MMLU*, and the *Aya Dataset*
- **MKC⁺:** Includes *MKC* and the *Aya Collection*

This allows us to evaluate the trade-off between data quality and scale.

Dataset Creation. For our model-based filtering approaches, our goal is to identify documents from the pretraining dataset that are most similar to our representative samples, with the notion of similarity determined by the specific classifier used. We can measure the similarity to our training dataset directly, for example, by computing the cosine similarity to our training samples in the embedding space. Alternatively, following the approach of Li et al. (2024b), the task can be framed as a binary classification problem, with the representative samples as the positive class. For the negative class, we can simply subsample documents from our pretraining dataset, under the assumption that the majority of these documents are neither well-structured nor knowledge-rich. We use both approaches for our classifiers.

To create the binary classification training dataset, we selected 80K random examples from the training set (*MKC* or *MKC⁺*) as positive samples and 80K random examples from FineWeb-2 as negative samples. For smaller datasets, such as *Include-Base-44*, the entire dataset was used. The same training dataset was utilized across all model-based filtering approaches, disregarding negative samples when unnecessary. Additionally, we created a training dataset for each language individually, to avoid leaking language-specific biases to data of other languages.

Sample Pre-processing. We applied no pre-processing to the FineWeb-2 (negative) samples but performed minimal pre-processing on the representative (positive) samples. For instance, in datasets like *MMLU* or *OpenAssistant-2*, we concatenated various sample components. For the *Aya Collection*, we resolved encoding issues in non-Latin languages and removed samples containing `<unk>` tokens, which were particularly prevalent in Arabic data (37.1%).

3.2 FASTTEXT-BASED FILTERING (FT)

To efficiently process datasets with over 100 million documents (Penedo et al., 2024c), similar to DCLM (Li et al., 2024b), we used a binary FastText classifier (Joulin et al., 2017). This classifier runs on CPU and can be easily deployed across multiple cores, for example using DataTrove (Penedo et al., 2024b).

We trained our FastText classifier on the processed training set using a 2-gram setting (4-gram for Chinese). Additional details about the training process are provided in Appendix A.1. These classifiers were then used to assign scores to all documents in the pretraining dataset. To filter the

dataset, we applied a score threshold based on the desired retention percentage of documents. This approach balances dataset size and the predicted quality of the samples.

3.3 TRANSFORMER EMBEDDING-BASED FILTERING

To leverage rich semantic information based on contextual relationships, we utilized the Transformer model embeddings. Specifically, we selected a pretrained XLM-RoBERTa base model (Conneau et al., 2020) due to its support of 100 languages, a relatively small size of approximately 279M parameters, and its transparent training procedure. This choice enabled us to process web-scale data efficiently without being restricted to a single language and to align with our commitment to open science.

To retain general embeddings that can be reused across methods, we opted against fine-tuning the model. For each document from our datasets, we computed the 768-dimensional embedding by mean pooling the embeddings of the output sequence. Since the model has a fixed maximum sequence length of 512 tokens, we considered only the first 512 tokens of each document, assuming they are representative of the entire document.

After computing the embeddings of our corpora, we experimented with two methods: 1) classification of embeddings using a multi-layer perceptron and 2) cosine similarity between the embeddings. As in the FastText approach, we scored each document and applied a threshold to retain the desired percentage of the highest-scoring documents.

Multi-Layer Perceptron (MLP). We trained a single-hidden-layer neural network with a hidden dimension of 256, the ReLU activation function, a dropout rate of 20%, and the sigmoid function on the output. The network was trained for 6 epochs using the AdamW optimizer (Loshchilov & Hutter, 2019) with a constant learning rate 0.0003 and binary cross-entropy loss. We computed document scores using the output layer of the MLP model, which used XLM-RoBERTa document embeddings as input.

Cosine Similarity (CS). We computed the document scores as the maximum cosine similarity between its embeddings and a set of K randomly sampled positive sample embeddings. We experimented with varying values of K , including 1024, 2048, 4096, 8192, and 16384. However, we did not observe a significant differences in the documents with high scores across these variations when manually inspecting the data. To strike a balance between the diversity of the positive samples and computational efficiency, we chose $K = 8192$ for our experiments.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

Technical Details. We evaluate 1B-parameter Llama models (Llama Team, 2024) to demonstrate the effectiveness of our model-based filtering approaches. The models are trained on either 70B or 119B tokens, balancing token quality and diversity. The smaller dataset (70B tokens) exposes the model to each token at most once (with a few exceptions where some tokens appear twice). The larger dataset (119B tokens) simulates longer training, resulting in increased token repetition. Training utilizes the HuggingFace Nanotron library (Hugging Face, 2024a) with the AdamW optimizer (Loshchilov & Hutter, 2019) and a WSD learning rate schedule (Hägele et al., 2024).

To minimize the need for costly hyperparameter tuning, we maintain a consistent setup across all experiments. Specifically, we adopt the DeepSeek scaling law (DeepSeek-AI et al., 2024) with a batch size of 1.6M tokens, learning rate of 0.0008, and 2000 warmup steps. We provide our Nanotron config in Appendix A.2.

As base dataset, we use FineWeb-2 (Penedo et al., 2024c), which has been shown to provide a strong baseline across a variety of languages. Since FineWeb-2 is globally deduplicated, we rehydrate both filtered and unfiltered data using the hyperparameters recommended by Penedo et al. (2024c).

To validate our method on English, we use three datasets: FineWeb (Penedo et al., 2024a) as the baseline, along with FineWeb-Edu (Penedo et al., 2024a) and DCLM (Li et al., 2024b), both of which represent the current state-of-the-art. Tokenization is performed using the multilingual Mistral v3

(Tekken) tokenizer (Mistral AI, 2024). All experiments are conducted using 80 NVIDIA GH200 chips.

Evaluation. Our evaluation prioritizes a diverse range of tasks to ensure the models retain well-rounded capabilities, rather than focusing exclusively on knowledge-based tasks. Specifically, we include tasks covering reading comprehension, general knowledge, natural language understanding, common-sense reasoning, and generative tasks in the target language. To evaluate our approach, we use the HuggingFace LightEval library (Fourrier et al., 2023).

For French, Chinese, and Arabic, we utilize the FineTasks (Kydlíček et al., 2024) multilingual evaluation suite, which is designed to provide meaningful signals even for models trained in the order of 100B tokens. We select analogous tasks for German and Danish. For English, we rely on the SmolLM tasks suite (Hugging Face, 2024b). A complete list of tasks and their evaluation metrics for each language is provided in Appendix C.

Model Selection. We follow the approach used in FineTasks for filter selection, computing a global rank score across individual metrics and languages to determine the optimal approach.

4.2 EXPERIMENTAL RESULTS & DISCUSSION

4.2.1 MODEL SELECTION

In Section 3, we introduced several model-based filtering approaches. *But which of these performs the best?* We evaluate which combination of our defined classifier training datasets (MKC or MKC^+) and filtering methods (FT , MLP or CS) achieve the highest performance. Table 1 presents the overall ranking across our representative language selection (Chinese, German, French, Arabic, Danish) and training runs of 70B and 119B tokens. Analogous to the DCLM filtering recipe (Li et al., 2024b), the results are based on a dataset that retains 10% of the documents for the high-resource datasets (Chinese, German, French) and keeps 56% and 65% of the documents for the lower-resource languages (Arabic and Danish, respectively). These percentages maintain approximately 70B tokens, under the assumption of uniform token distribution across documents. We also exclude approaches that use MKC for training on Danish, as it lacks sufficient training data. For detailed, per-language results, please refer to Appendix B.1.

Table 1 demonstrates that $MLP\ MKC^+$ approach outperforms all other approaches. Interestingly, the high- and low-scored samples presented in Appendix D align with the observed rankings. Figure 1 further highlights the strong performance of $MLP\ MKC^+$, particularly for high-resource languages, where it largely outperforms the baseline. For lower-resource languages—where less data was filtered—the performance gains are less pronounced. Notably, FT filtering is also competitive. Given the computational expense of XLM-RoBERTa embeddings, FastText can be a promising alternative in resource-constrained setups.

Table 1: Benchmark performance comparison (average rank) between the baseline (FineWeb-2) and our proposed filtering methods (FT , MLP , and CS) trained on MKC^+ or MKC , retaining top 10% of the documents for Chinese, German, and French, 56% for Arabic, and 65% for Danish. The average rank is computed across FineTasks performance of 1B-parameter models evaluated after 70B and 119B tokens were consumed.

Approach	Average Rank
$MLP\ MKC^+$	4.35
$MLP\ MKC$	6.11
$FT\ MKC^+$	7.17
$FT\ MKC$	8.04
$CS\ MKC$	8.10
Baseline	8.72
$CS\ MKC^+$	8.79

4.2.2 THRESHOLD SELECTION

In Section 4.2.1, we base our model selection on experiments that retain top 10% of the data for high-resource languages. *But is this the optimal threshold?* Following the methodology of Li et al. (2024b),




Figure 1: Benchmark performance comparison (accuracy) during training between the baseline methods (FineWeb, DCLM, FineWeb-Edu, and FineWeb-2) and our proposed filtering methods (*FT*, *MLP*, and *CS*), trained on *MKC*⁺. When using our approaches, the data retention rates are set to 10% for English, Chinese, German, and French, 56% for Arabic, and 65% for Danish.

we analyze the impact of varying filter strengths on performance for Chinese, German, and French, using our *MLP* and *FT* filtering methods. The results are summarized in Table 2, with a comprehensive analysis, including results for *CS*, provided in Appendix B.2 (Table 14). Consistent with their findings, we observe that retaining top 10% of the data is a competitive threshold, particularly for approaches using the *MKC*⁺ dataset. Interestingly, approaches using *MKC* perform better with higher retention. Motivated by the observed bias in certain approaches favoring the selection of shorter documents, we examine how this bias interacts with performance when retaining more documents. As demonstrated in Figure 2 for German, Appendix B.2 for other languages, and the retained token counts in Table 15, the *MLP MKC* approach shows a tendency to retain shorter documents, while achieving higher performance with an increased number of retained documents. In contrast, the *CS* and *FT* filtering methods present mixed results, suggesting that the optimal threshold selection may be influenced by additional factors.

Table 2: Benchmark performance comparison (average rank) between the baseline (FineWeb-2) and our proposed filtering methods (*FT*, *MLP*) trained on *MKC*⁺ or *MKC*, retaining top 10%, 15% or 20% of the documents. The average rank is computed across FineTasks performance of 1B-parameter models evaluated for Chinese, German and French after 70B and 119B tokens were consumed.

Approach	Threshold	Average Rank
<i>MLP MKC</i> ⁺	10%	8.85
<i>MLP MKC</i> ⁺	15%	9.44
<i>MLP MKC</i>	20%	11.37
<i>MLP MKC</i>	15%	11.70
<i>MLP MKC</i>	10%	11.95
<i>MLP MKC</i> ⁺	20%	11.97
<i>FT MKC</i> ⁺	10%	13.92
<i>FT MKC</i>	15%	14.62
<i>FT MKC</i>	10%	14.74
<i>FT MKC</i>	20%	15.62
<i>FT MKC</i> ⁺	15%	16.27
<i>FT MKC</i> ⁺	20%	16.51
Baseline	–	18.55

4.2.3 TRAINING DATA ANALYSIS

The experiments in Sections 4.2.1 and 4.2.2 are based on the training datasets *MKC* and *MKC*⁺. *But is the diversity introduced by combining various base datasets truly necessary?* We evaluate the




Figure 2: Comparison of average document length and standard deviation in FineWeb-2 before and after filtering using one of our approaches retaining top 10% of the documents. The average document length of FineWeb-2 is represented as a red horizontal line, while the medians are shown as red dots. Document length is measured based on number of space-separated tokens.

impact of each base dataset individually and compare it to the combined MKC^+ dataset. For this ablation study, we use our best filtering method (MLP with a top 10% retention) and train the models on 30B tokens. This token count is chosen to match the size of the smallest filtered dataset, ensuring consistency across comparisons. The results, presented in Table 3, show that despite the absence of a quality guarantee for all samples in the *Aya Collection*, this dataset yields strong performance, making our approach applicable for various languages. Overall, we observe that the diversity resulting from combining all individual training datasets gives the best results.

Interestingly, models trained exclusively on *Include-Base-44* and *OpenAssistant-2* perform worse overall than the baseline. This may be due to the nature of these datasets. For instance, *Include-Base-44* is relatively small and domain-specific, e.g., consisting primarily of driving license exam questions in its German subset. Similarly, *OpenAssistant-2* includes a limited number of samples, with fewer than 2K positive samples per training set, which likely negatively impacts classifier performance. Again, we relate model performance to the average document length bias in Appendix B.3 and confirm the findings from Section 4.2.2, suggesting that factors beyond the retained document length bias may influence performance.

Table 3: Benchmark performance comparison (average rank) between the baseline (FineWeb-2) and the MLP filtering method trained on either MKC^+ as a whole or its individual dataset components, retaining top 10% of the documents for Chinese, German, and French, 56% for Arabic, and 65% for Danish. The average rank is computed across FineTasks performance of 1B-parameter models trained on each language with 30B tokens.

Dataset	Average Rank
MKC^+	2.52
<i>Aya Collection</i>	2.91
<i>Aya Dataset</i>	3.17
<i>MMLU</i>	3.57
Baseline	4.09
<i>OpenAssistant-2</i>	4.53
<i>Include-Base-44</i>	5.42

Table 4: Benchmark performance comparison (average rank) of our $MLP MKC^+$ and $FT MKC^+$ approaches, retaining top 10% of the documents while mixing in 0%, 5% or 10% of the FineWeb-2 dataset. The average rank is computed across FineTasks performance of 1B-parameter models evaluated for Chinese, German, or French, after consuming 70B and 119B tokens.

Approach	Mixture Rate	Average Rank
$MLP MKC^+$	5%	5.09
$MLP MKC^+$	0%	5.16
$MLP MKC^+$	10%	5.40
$FT MKC^+$	10%	7.17
$FT MKC^+$	0%	7.51
$FT MKC^+$	5%	8.66

4.2.4 DATASET MIXING

But does our model-based filtering retain data diversity? We explore whether incorporating a small percentage of raw data can help improve performance. We do this for our best FastText ($FT MKC^+$) and Transformer approaches ($MLP MKC^+$). Table 4 presents the results of experiments where 5% and 10% unfiltered data were mixed into the training dataset, alongside results from training without any data mixing. Both $FT MKC^+$ and $MLP MKC^+$ approaches show mixed signal, although $MLP MKC^+$ approach demonstrates little difference between mixing 5% unfiltered data and no mixing, indicating that it retains enough diversity.

4.2.5 APPROACH VALIDATION ON ENGLISH

Previous experiments have shown strong performance of our *MLP MKC⁺* approach. *But do these results translate to English?* Table 5 presents the performance of *MLP MKC⁺* with 10% retention applied to the English FineWeb dataset Penedo et al. (2024a). Our method is compared against FineWeb and baselines using model-based filtered datasets, including DCLM Li et al. (2024b) and FineWeb-Edu (Penedo et al., 2024a). To save computational resources, we use the 6 most recent FineWeb and FineWeb-Edu dumps and the first partition of DCLM⁵ (Penedo et al., 2024a), which we denote with *. Each of these subsets contains more than 119B tokens, with FineWeb retaining this size even after applying our filtering retaining top 10% of the documents.

While each approach demonstrates strengths in different benchmarks, as seen from Table 5 and Figure 1, the overall average rank results indicate that our method outperforms all other baselines.

Table 5: Benchmark performance comparison for English of our *MLP MKC⁺* approach (retaining top 10% of the documents) against baseline datasets: FineWeb, DCLM, and FineWeb-Edu. The average rank is computed across SmolLM task performance for 1B-parameter models trained on 119B tokens.

Dataset	Ours	DCLM*	FW-Edu*	FW*
Average Rank	1.8333	2.3889	2.4444	3.3333
ARC (Challenge)	0.3550	0.3530	0.3850	0.3010
ARC (Easy)	0.6670	0.6470	0.6970	0.5880
CommonsenseQA	0.3870	0.4100	0.3770	0.3850
HellaSwag	0.6040	0.5960	0.5700	0.5930
MMLU	0.3400	0.3160	0.3470	0.3030
OpenBookQA	0.3860	0.3840	0.4180	0.3560
PIQA	0.7510	0.7510	0.7410	0.7620
WinoGrande	0.5720	0.5610	0.5660	0.5550
TriviaQA	0.0820	0.1240	0.0320	0.0370

4.2.6 DATA CONTAMINATION ANALYSIS

Our LLMs are never trained on benchmark datasets. *But is the strong performance observed in the previous sections primarily due to an increased ratio of data contamination?* To ensure the validity of our approach, we conduct decontamination experiments, as web crawl data may include evaluation benchmark tasks. While Li et al. (2024b) addressed similar concerns, our approach follows the methodology of Brown et al. (2020). Specifically, we perform 13-gram decontamination of the LLM training data separately for English and French evaluation benchmarks. However, unlike the original approach, we remove the entire document if it is flagged as contaminated, using the implementation provided in DataTrove (Penedo et al., 2024b).

Tables 6 and 7 present the results of decontamination experiments for English and French, respectively. We used the following experimental setup (removed document contamination rates): baseline FineWeb English (0.16%), *MLP MKC⁺* English with 10% retention (0.19%), baseline FineWeb-2 French (0.14%), and *MLP MKC⁺* French with 10% retention (0.14%). As in our previous experiments, we train the models on 119B tokens. Additionally, we compare the results against equivalent training runs without decontamination to further analyze its impact. For an example of a contaminated sample, see Appendix E.

For English models, decontamination slightly reduces performance both for our approach and baseline FineWeb data. However, even when decontaminated, our approach still outperforms training on non-decontaminated baseline data. For French models, performance of our approach is comparable between decontaminated and non-decontaminated datasets, with both continuing to outperform baseline FineWeb-2 data. Interestingly, decontaminated baseline data yields better results than its non-decontaminated counterpart.

4.2.7 IMPACT ON MULTILINGUAL MODEL TRAINING

Although not the primary focus of our work, we believe that refined datasets can contribute to advancing the performance of multilingual models. To investigate this, we conducted an ablation

⁵huggingface.co/datasets/mlfoundations/dclm-baseline-1.0-parquet

Table 6: Benchmark performance comparison in English for our *MLP MKC*⁺ approach (retaining top 10% of the documents), both decontaminated (*D*) and non-decontaminated, against the baseline FineWeb datasets, also in decontaminated and non-decontaminated variants. The average rank is computed across SmoLLM task performance for 1B-parameter models trained on 119B tokens.

Dataset	Ours	Ours _D	FW*	FW _D *
Average Rank	1.5000	2.1111	3.0556	3.3333
ARC (Challenge)	0.3550	0.3440	0.3010	0.2880
ARC (Easy)	0.6670	0.6520	0.5880	0.5700
CommonsenseQA	0.3870	0.4000	0.3850	0.3820
HellaSwag	0.6040	0.6040	0.5930	0.5890
MMLU	0.3400	0.3220	0.3030	0.3050
OpenBookQA	0.3860	0.3840	0.3560	0.3740
PIQA	0.7510	0.7590	0.7620	0.7600
WinoGrande	0.5720	0.5550	0.5550	0.5570
TriviaQA	0.0820	0.0380	0.0370	0.0250

Table 7: Benchmark performance comparison in French for our *MLP MKC*⁺ approach (retaining top 10% of the documents), both decontaminated (*D*) and non-decontaminated, against the baseline FineWeb-2 datasets, also in decontaminated and non-decontaminated variants. The average rank is computed across FineTasks performance for 1B-parameter models trained on 119B tokens.

Dataset	Ours	Ours _D	FW-2 _D	FW-2
Average Rank	2.0556	2.0556	2.7222	3.1667
Belebele	0.3533	0.3400	0.3778	0.3444
HellaSwag	0.5380	0.5350	0.5180	0.5180
X-CSQA	0.2740	0.2810	0.2730	0.2870
XNLI 2.0	0.7400	0.7400	0.7070	0.7180
FQuAD	0.2803	0.2620	0.2890	0.2401
MMLU	0.2895	0.2875	0.2711	0.2706
Mintaka	0.0438	0.0797	0.0658	0.0712
X-CODAH	0.2667	0.2900	0.2800	0.2633
ARC (Challenge)	0.3180	0.3110	0.2880	0.2850

study by training a 1B-parameter model on 595B tokens ($5 \times 119B$), covering all five languages: Chinese, German, French, Arabic and Danish. We trained two models—the first one using our filtered FineWeb-2 dataset and the second one using unfiltered FineWeb-2 data. We then compared these results for each language against their monolingual counterparts trained on 119B tokens.

The results for French are presented in Table 8. We observe that the multilingual LLM outperforms its monolingual counterpart on our filtered datasets, whereas the monolingual model achieves better performance than the multilingual model on the FineWeb-2 dataset. This trend is consistent across all languages except Chinese. Detailed results for the other languages are provided in Appendix B.4.

Table 8: Benchmark performance comparison for French of multilingual LLMs (*M*) trained on FineWeb-2 or the refined dataset using our *MLP MKC*⁺ approach (retaining top 10% of the documents for Chinese, German, and French, 56% for Arabic, and 65% for Danish) trained on 595B tokens, against their monolingual counterparts trained on 119B tokens. The average rank is computed across FineTasks performance for 1B-parameter models trained on 119B tokens.

Dataset	Ours _M	Ours	FW-2	FW-2 _M
Average Rank	1.8333	2.0556	3.0000	3.1111
Belebele	0.3667	0.3533	0.3444	0.3511
HellaSwag	0.5270	0.5380	0.5180	0.4970
X-CSQA	0.2740	0.2740	0.2870	0.2750
XNLI 2.0	0.7660	0.7400	0.7180	0.7330
FQuAD	0.3212	0.2803	0.2401	0.2459
MMLU	0.2841	0.2895	0.2706	0.2735
Mintaka	0.0456	0.0438	0.0712	0.0579
X-CODAH	0.2900	0.2667	0.2633	0.2567
ARC (Challenge)	0.2970	0.3180	0.2850	0.2670

5 CONCLUSION

In this work, we introduced a novel framework for model-based filtering of web-scale multilingual pretraining datasets, demonstrating consistent improvements on LLM benchmarks across a wide range of languages. Our Transformer embedding-based classifier, *MLP MKC*⁺, outperforms state-of-the-art methods on both English and multilingual datasets, even when decontaminating the datasets or using them for training multilingual LLMs. This demonstrates that simple classifiers can achieve competitive results. While our FastText-based filtering approach performed well and shows promise in resource-constrained setups, *MLP MKC*⁺ consistently outperformed all other methods and can be easily scaled to other languages. These results motivate us to expand our framework to 20 languages and release the corresponding refined pretraining datasets and our code, contributing to the advancement of multilingual language modeling.

REFERENCES

- Julien Abadji, Pedro Javier Ortiz Suárez, Laurent Romary, and Benoît Sagot. Ungoliant: An optimized pipeline for the generation of a very large-scale multilingual web corpus. *Proceedings of the Workshop on Challenges in the Management of Large Corpora (CMLC-9) 2021*. Limerick, 12 July 2021 (Online-Event), pp. 1 – 9, Mannheim, 2021. Leibniz-Institut für Deutsche Sprache. doi: 10.14618/ids-pub-10468. URL <https://nbn-resolving.org/urn:nbn:de:bsz:mh39-104688>.
- Julien Abadji, Pedro Ortiz Suarez, Laurent Romary, and Benoît Sagot. Towards a Cleaner Document-Oriented Multilingual Crawled Corpus. *arXiv e-prints*, art. arXiv:2201.06642, January 2022.
- Ebtesam Almazrouei, Ruxandra Cojocaru, Michele Baldo, Quentin Malartic, Hamza Alobeidli, Daniele Mazzotta, Guilherme Penedo, Giulia Campesan, Mugariya Farooq, Maitha Alhammedi, Julien Launay, and Badreddine Noune. AlGhafa evaluation benchmark for Arabic language models. In Hassan Sawaf, Samhaa El-Beltagy, Wajdi Zaghouani, Walid Magdy, Ahmed Abdelali, Nadi Tomeh, Ibrahim Abu Farha, Nizar Habash, Salam Khalifa, Amr Keleg, Hatem Haddad, Imed Zitouni, Khalil Mrini, and Rawan Almatham (eds.), *Proceedings of ArabicNLP 2023*, pp. 244–275, Singapore (Hybrid), December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.arabnlp-1.21. URL <https://aclanthology.org/2023.arabnlp-1.21/>.
- Zachary Ankner, Cody Blakeney, Kartik Sreenivasan, Max Marion, Matthew L. Leavitt, and Mansheej Paul. Perplexed by perplexity: Perplexity-based data pruning with small reference models, 2024. URL <https://arxiv.org/abs/2405.20541>.
- Catherine Arnett, Eliot Jones, Ivan P. Yamshchikov, and Pierre-Carl Langlais. Toxicity of the Commons: Curating Open-Source Pre-Training Data. *arXiv preprint arXiv:2410.22587*, 2024. URL <https://arxiv.org/pdf/2410.22587>.
- Lucas Bandarkar, Davis Liang, Benjamin Muller, Mikel Artetxe, Satya Narayan Shukla, Donald Husa, Naman Goyal, Abhinandan Krishnan, Luke Zettlemoyer, and Madian Khabsa. The Belebele Benchmark: a Parallel Reading Comprehension Dataset in 122 Language Variants. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 749–775. Association for Computational Linguistics, 2024. doi: 10.18653/v1/2024.acl-long.44. URL <http://dx.doi.org/10.18653/v1/2024.acl-long.44>.
- Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. Piqa: Reasoning about physical commonsense in natural language, 2019. URL <https://arxiv.org/abs/1911.11641>.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information, 2017. URL <https://arxiv.org/abs/1607.04606>.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. TyDi QA: A Benchmark for Information-Seeking Question Answering in Typologically Diverse Languages, 2020. URL <https://arxiv.org/abs/2003.05002>.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have Solved Question Answering? Try ARC, the AI2 Reasoning Challenge, 2018. URL <https://arxiv.org/abs/1803.05457>.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. XNLI: Evaluating cross-lingual sentence representations. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun’ichi Tsujii (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2475–2485, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1269. URL <https://aclanthology.org/D18-1269/>.

- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. Unsupervised cross-lingual representation learning at scale, 2020. URL <https://arxiv.org/abs/1911.02116>.
- Yiming Cui, Ting Liu, Wanxiang Che, Li Xiao, Zhipeng Chen, Wentao Ma, Shijin Wang, and Guoping Hu. A span-extraction dataset for chinese machine reading comprehension. 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)*. Association for Computational Linguistics, 2019. doi: 10.18653/v1/d19-1600. URL <http://dx.doi.org/10.18653/v1/D19-1600>.
- Ona de Gibert, Graeme Nail, Nikolay Arefyev, Marta Bañón, Jelmer van der Linde, Shaoxiong Ji, Jaume Zaragoza-Bernabeu, Mikko Aulamo, Gema Ramírez-Sánchez, Andrey Kutuzov, Sampo Pyysalo, Stephan Open, and Jörg Tiedemann. A new massive multilingual dataset for high-performance language technologies. In Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (eds.), *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pp. 1116–1128, Torino, Italia, May 2024. ELRA and ICCL. URL <https://aclanthology.org/2024.lrec-main.100>.
- Ona De Gibert, Graeme Nail, Nikolay Arefyev, Marta Bañón, Jelmer Van Der Linde, Shaoxiong Ji, Jaume Zaragoza-Bernabeu, Mikko Aulamo, Gema Ramírez-Sánchez, Andrey Kutuzov, et al. A new massive multilingual dataset for high-performance language technologies. *arXiv preprint arXiv:2403.14009*, 2024.
- 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. DeepSeek LLM: Scaling Open-Source Language Models with Longtermism, 2024. URL <https://arxiv.org/abs/2401.02954>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2019. URL <https://arxiv.org/abs/1810.04805>.
- Martin d’Hoffschmidt, Wacim Belblidia, Tom Brendlé, Quentin Heinrich, and Maxime Vidal. FQuAD: French Question Answering Dataset, 2020. URL <https://arxiv.org/abs/2002.06071>.
- Sophie Fischer, Federico Rossetto, Carlos Gemmell, Andrew Ramsay, Iain Mackie, Philip Zobel, Niklas Tecklenburg, and Jeffrey Dalton. Open assistant toolkit–version 2. *arXiv preprint arXiv:2403.00586*, 2024.
- Clémentine Fourier, Nathan Habib, Thomas Wolf, and Lewis Tunstall. LightEval: A lightweight framework for LLM evaluation, 2023. URL <https://github.com/huggingface/lighteval>.
- 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. Textbooks are all you need, 2023. URL <https://arxiv.org/abs/2306.11644>.

- Alexander Hägele, Elie Bakouch, Atli Kosson, Loubna Ben Allal, Leandro Von Werra, and Martin Jaggi. Scaling laws and compute-optimal training beyond fixed training durations. *arXiv preprint arXiv:2405.18392*, 2024.
- Momchil Hardalov, Todor Mihaylov, Dimitrina Zlatkova, Yoan Dinkov, Ivan Koychev, and Preslav Nakov. Exams: A multi-subject high school examinations dataset for cross-lingual and multilingual question answering, 2020. URL <https://arxiv.org/abs/2011.03080>.
- William Held, Bhargavi Paranjape, Punit Singh Koura, Mike Lewis, Frank Zhang, and Todor Mihaylov. Optimizing pretraining data mixtures with llm-estimated utility. *arXiv preprint arXiv:2501.11747*, 2025.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020.
- Hai Hu, Kyle Richardson, Liang Xu, Lu Li, Sandra Kübler, and Lawrence Moss. OCNLI: Original Chinese Natural Language Inference. In Trevor Cohn, Yulan He, and Yang Liu (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 3512–3526, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.314. URL <https://aclanthology.org/2020.findings-emnlp.314/>.
- Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, Junlei Zhang, Jinghan Zhang, Tangjun Su, Junteng Liu, Chuancheng Lv, Yikai Zhang, Jiayi Lei, Yao Fu, Maosong Sun, and Junxian He. C-eval: A multi-level multi-discipline chinese evaluation suite for foundation models, 2023. URL <https://arxiv.org/abs/2305.08322>.
- Hugging Face. Nanotron, 2024a. URL <https://github.com/huggingface/nanotron>. Accessed 30 Jan. 2025.
- Hugging Face. SmolLM - blazingly fast and remarkably powerful, 2024b. URL <https://huggingface.co/blog/smollm>. Accessed 30 Jan. 2025.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In Regina Barzilay and Min-Yen Kan (eds.), *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1601–1611, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1147. URL <https://aclanthology.org/P17-1147/>.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pp. 427–431. Association for Computational Linguistics, April 2017.
- Fajri Koto, Haonan Li, Sara Shatnawi, Jad Doughman, Abdelrahman Boda Sadallah, Aisha Alraeesi, Khalid Almubarak, Zaid Alyafeai, Neha Sengupta, Shady Shehata, Nizar Habash, Preslav Nakov, and Timothy Baldwin. Arabicmmlu: Assessing massive multitask language understanding in arabic, 2024. URL <https://arxiv.org/abs/2402.12840>.
- Sneha Kudugunta, Isaac Caswell, Biao Zhang, Xavier Garcia, Christopher A. Choquette-Choo, Katherine Lee, Derrick Xin, Aditya Kusupati, Romi Stella, Ankur Bapna, and Orhan Firat. MADLAD-400: A Multilingual And Document-Level Large Audited Dataset, 2023. URL <https://arxiv.org/abs/2309.04662>.
- Hynek Kydlíček, Guilherme Penedo, Clémentine Fourier, Nathan Habib, and Thomas Wolf. FineTasks: Finding signal in a haystack of 200+ multilingual tasks, 2024. URL <https://huggingface.co/spaces/HuggingFaceFW/blogpost-fine-tasks>. Accessed 30 Jan. 2025.
- Viet Lai, Chien Nguyen, Nghia Ngo, Thuat Nguyen, Franck Dernoncourt, Ryan Rossi, and Thien Nguyen. Okapi: Instruction-tuned large language models in multiple languages with reinforcement learning from human feedback. In Yansong Feng and Els Lefever (eds.), *Proceedings of*

- the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 318–327, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-demo.28. URL <https://aclanthology.org/2023.emnlp-demo.28/>.
- Guillaume Lample and Alexis Conneau. Cross-lingual language model pretraining, 2019. URL <https://arxiv.org/abs/1901.07291>.
- Hugo Laurençon, Lucile Saulnier, Thomas Wang, Christopher Akiki, Albert Villanova del Moral, Teven Le Scao, Leandro Von Werra, Chenghao Mou, Eduardo González Ponferrada, Huu Nguyen, et al. The bigscience roots corpus: A 1.6 tb composite multilingual dataset. *Advances in Neural Information Processing Systems*, 35:31809–31826, 2022.
- Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. Deduplicating training data makes language models better. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 8424–8445, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.577. URL <https://aclanthology.org/2022.acl-long.577/>.
- Patrick Lewis, Barlas Ögüz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. Mlqa: Evaluating cross-lingual extractive question answering, 2020. URL <https://arxiv.org/abs/1910.07475>.
- Haonan Li, Yixuan Zhang, Fajri Koto, Yifei Yang, Hai Zhao, Yeyun Gong, Nan Duan, and Timothy Baldwin. Cmmlu: Measuring massive multitask language understanding in chinese, 2024a. URL <https://arxiv.org/abs/2306.09212>.
- Jeffrey Li, Alex Fang, Georgios Smyrnis, Maor Ivgi, Matt Jordan, Samir Gadre, Hritik Bansal, Etash Guha, Sedrick Keh, Kushal Arora, et al. DataComp-LM: In search of the next generation of training sets for language models. *arXiv preprint arXiv:2406.11794*, 2024b.
- Bill Yuchen Lin, Seyeon Lee, Xiaoyang Qiao, and Xiang Ren. Common sense beyond English: Evaluating and improving multilingual language models for commonsense reasoning. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1274–1287, Online, August 2021a. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.102. URL <https://aclanthology.org/2021.acl-long.102/>.
- Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, Daniel Simig, Myle Ott, Naman Goyal, Shruti Bhosale, Jingfei Du, Ramakanth Pasunuru, Sam Shleifer, Punit Singh Koura, Vishrav Chaudhary, Brian O’Horo, Jeff Wang, Luke Zettlemoyer, Zornitsa Kozareva, Mona T. Diab, Veselin Stoyanov, and Xian Li. Few-shot learning with multilingual language models. *CoRR*, abs/2112.10668, 2021b. URL <https://arxiv.org/abs/2112.10668>.
- Llama Team. The Llama 3 Herd of Models, 2024. URL <https://arxiv.org/abs/2407.21783>.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. 2019. URL <https://arxiv.org/abs/1711.05101>.
- Max Marion, Ahmet Üstün, Luiza Pozzobon, Alex Wang, Marzieh Fadaee, and Sara Hooker. When less is more: Investigating data pruning for pretraining llms at scale, 2023. URL <https://arxiv.org/abs/2309.04564>.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering, 2018. URL <https://arxiv.org/abs/1809.02789>.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space, 2013. URL <https://arxiv.org/abs/1301.3781>.

- Mistral AI. v3 (tekken) tokenizer, 2024. URL <https://docs.mistral.ai/guides/tokenization/>. Accessed 30 Jan. 2025.
- Mistral AI. Mistral small 3, 2025. URL <https://mistral.ai/news/mistral-small-3/>. Accessed 30 Jan. 2025.
- Hussein Mozannar, Karl El Hajal, Elie Maamary, and Hazem Hajj. Neural arabic question answering, 2019. URL <https://arxiv.org/abs/1906.05394>.
- Niklas Muennighoff, Alexander M. Rush, Boaz Barak, Teven Le Scao, Aleksandra Piktus, Nouamane Tazi, Sampo Pyysalo, Thomas Wolf, and Colin Raffel. Scaling data-constrained language models, 2023. URL <https://arxiv.org/abs/2305.16264>.
- Thuat Nguyen, Chien Van Nguyen, Viet Dac Lai, Hieu Man, Nghia Trung Ngo, Franck Dernoncourt, Ryan A Rossi, and Thien Huu Nguyen. Culturax: A cleaned, enormous, and multilingual dataset for large language models in 167 languages. *arXiv preprint arXiv:2309.09400*, 2023.
- OpenAI. MMMLU, 2024. URL <https://huggingface.co/datasets/openai/MMMLU>. Accessed 30 Jan. 2025.
- Pedro Javier Ortiz Suárez, Benoît Sagot, and Laurent Romary. Asynchronous pipelines for processing huge corpora on medium to low resource infrastructures. *Proceedings of the Workshop on Challenges in the Management of Large Corpora (CMLC-7) 2019*. Cardiff, 22nd July 2019, pp. 9–16, Mannheim, 2019. Leibniz-Institut für Deutsche Sprache. doi: 10.14618/ids-pub-9021. URL <http://nbn-resolving.de/urn:nbn:de:bsz:mh39-90215>.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Hamza Alobeidli, Alessandro Cappelli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. The RefinedWeb dataset for Falcon LLM: Outperforming curated corpora with web data only. *Advances in Neural Information Processing Systems*, 36:79155–79172, 2023.
- Guilherme Penedo, Hynek Kydlíček, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, Thomas Wolf, et al. The FineWeb Datasets: Decanting the Web for the Finest Text Data at Scale. *arXiv preprint arXiv:2406.17557*, 2024a.
- Guilherme Penedo, Hynek Kydlíček, Alessandro Cappelli, Mario Sasko, and Thomas Wolf. Data-Trove: large scale data processing, 2024b. URL <https://github.com/huggingface/datatrove>. Accessed 30 Jan. 2025.
- Guilherme Penedo, Hynek Kydlíček, Vinko Sabolčec, Bettina Messmer, Negar Foroutan, Martin Jaggi, Leandro von Werra, and Thomas Wolf. FineWeb2: A sparkling update with 1000s of languages, December 2024c. URL <https://huggingface.co/datasets/HuggingFaceFW/fineweb-2>. Accessed 30 Jan. 2025.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. GloVe: Global vectors for word representation. In Alessandro Moschitti, Bo Pang, and Walter Daelemans (eds.), *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1532–1543, Doha, Qatar, October 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1162. URL <https://aclanthology.org/D14-1162/>.
- Pluto-Junzeng. pluto-junzeng/chinesesquad, 2019. URL <https://github.com/pluto-junzeng/ChineseSquad>. Accessed 30 Jan. 2025.
- Edoardo Maria Ponti, Goran Glavaš, Olga Majewska, Qianchu Liu, Ivan Vulić, and Anna Korhonen. XCOPA: A multilingual dataset for causal commonsense reasoning. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2362–2376, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.185. URL <https://aclanthology.org/2020.emnlp-main.185/>.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. 2018.

- Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*, 2021.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
- Angelika Romanou, Negar Foroutan, Anna Sotnikova, Zeming Chen, Sree Harsha Nelaturu, Shivalika Singh, Rishabh Maheshwary, Micol Altomare, Mohamed A Haggag, Alfonso Amayuelas, et al. Include: Evaluating multilingual language understanding with regional knowledge. *arXiv preprint arXiv:2411.19799*, 2024.
- Naveen Sachdeva, Benjamin Coleman, Wang-Cheng Kang, Jianmo Ni, Lichan Hong, Ed H. Chi, James Caverlee, Julian McAuley, and Derek Zhiyuan Cheng. How to train data-efficient llms, 2024. URL <https://arxiv.org/abs/2402.09668>.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. WinoGrande: An Adversarial Winograd Schema Challenge at Scale. *arXiv preprint arXiv:1907.10641*, 2019.
- Priyanka Sen, Alham Fikri Aji, and Amir Saffari. Mintaka: A Complex, Natural, and Multilingual Dataset for End-to-End Question Answering, 2022. URL <https://arxiv.org/abs/2210.01613>.
- Shivalika Singh, Angelika Romanou, Cl  mentine Fourrier, David I. Adelani, Jian Gang Ngui, Daniel Vila-Suero, Peerat Limkonchotiawat, Kelly Marchisio, Wei Qi Leong, Yosephine Susanto, Raymond Ng, Shayne Longpre, Wei-Yin Ko, Madeline Smith, Antoine Bosselut, Alice Oh, Andre F. T. Martins, Leshem Choshen, Daphne Ippolito, Enzo Ferrante, Marzieh Fadaee, Beyza Ermi  s, and Sara Hooker. Global mmlu: Understanding and addressing cultural and linguistic biases in multilingual evaluation, 2024a. URL <https://arxiv.org/abs/2412.03304>.
- Shivalika Singh, Freddie Vargus, Daniel Dsouza, B  rje 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 Krzemi  nski, Hakimeh Fadaei, Irem Erg  n, 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   st  n, Marzieh Fadaee, and Sara Hooker. Aya dataset: An open-access collection for multilingual instruction tuning, 2024b.
- Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Authur, Ben Bogin, Khyathi Chandu, Jennifer Dumas, Yanai Elazar, et al. Dolma: An open corpus of three trillion tokens for language model pretraining research. *arXiv preprint arXiv:2402.00159*, 2024.
- Nishant Subramani, Sasha Luccioni, Jesse Dodge, and Margaret Mitchell. Detecting personal information in training corpora: an analysis. In Anaelia Ovalle, Kai-Wei Chang, Ninareh Mehrabi, Yada Pruksachatkun, Aram Galystan, Jwala Dhamala, Apurv Verma, Trista Cao, Anoop Kumar, and Rahul Gupta (eds.), *Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing (TrustNLP 2023)*, pp. 208–220, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.trustnlp-1.18. URL <https://aclanthology.org/2023.trustnlp-1.18/>.
- Kai Sun, Dian Yu, Dong Yu, and Claire Cardie. Investigating prior knowledge for challenging Chinese machine reading comprehension. *Transactions of the Association for Computational Linguistics*, 8:141–155, 2020. doi: 10.1162/tacl.a.00305. URL <https://aclanthology.org/2020.tacl-1.10/>.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. CommonsenseQA: A question answering challenge targeting commonsense knowledge. In Jill Burstein, Christy Doran, and Tamar Solorio (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4149–4158, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1421. URL <https://aclanthology.org/N19-1421/>.

- Alexey Tikhonov and Max Ryabinin. It’s all in the heads: Using attention heads as a baseline for cross-lingual transfer in commonsense reasoning, 2021. URL <https://arxiv.org/abs/2106.12066>.
- Together Computer. Redpajama: An open source recipe to reproduce llama training dataset, 2023. URL <https://github.com/togethercomputer/RedPajama-Data>. Accessed 30 Jan. 2025.
- Ankit Kumar Upadhyay and Harsit Kumar Upadhyay. Xnli 2.0: Improving xnli dataset and performance on cross lingual understanding (xlu), 2023. URL <https://arxiv.org/abs/2301.06527>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023. URL <https://arxiv.org/abs/1706.03762>.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Edouard Grave. Ccnet: Extracting high quality monolingual datasets from web crawl data. *arXiv preprint arXiv:1911.00359*, 2019.
- Alexander Wettig, Aatmik Gupta, Saumya Malik, and Danqi Chen. Qurating: Selecting high-quality data for training language models, 2024. URL <https://arxiv.org/abs/2402.09739>.
- Sang Michael Xie, Shibani Santurkar, Tengyu Ma, and Percy Liang. Data selection for language models via importance resampling, 2023. URL <https://arxiv.org/abs/2302.03169>.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. mT5: A massively multilingual pre-trained text-to-text transformer. In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 483–498, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.41. URL <https://aclanthology.org/2021.naacl-main.41/>.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine really finish your sentence?, 2019. URL <https://arxiv.org/abs/1905.07830>.
- Wenxuan Zhang, Sharifah Mahani Aljunied, Chang Gao, Yew Ken Chia, and Lidong Bing. M3exam: A multilingual, multimodal, multilevel benchmark for examining large language models, 2023. URL <https://arxiv.org/abs/2306.05179>.
- Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. Agieval: A human-centric benchmark for evaluating foundation models, 2023. URL <https://arxiv.org/abs/2304.06364>.

A ADDITIONAL EXPERIMENTAL DETAILS

A.1 FASTTEXT TRAINING DETAILS

The FastText classifier was trained on the processed training set using 2-grams, a minCount of 1, and the softmax loss function. All other parameters were automatically tuned using the FastText library. For Chinese, fixed parameters were used: 30 training epochs and a learning rate of 0.1 to ensure training stability. Additionally, 4-grams and a minCount of 0 were selected based on manual evaluation of the results.

Prior to training the FastText models, we pre-processed the training data by removing newlines.

A.2 NANOTRON CONFIGURATION

To facilitate the reproducibility of our model training, we provide the Nanotron (Hugging Face, 2024a) configuration used in our experiments.

```

1 checkpoints:
2   checkpoint_interval: 1000
3   checkpoints_path: checkpoints/
4   checkpoints_path_is_shared_file_system: false
5   resume_checkpoint_path: null
6   save_initial_state: false
7 data_stages:
8   - data:
9       dataset:
10         dataset_folder: template
11         num_loading_workers: 1
12         seed: 42
13         name: General purpose training (Single dataset)
14         start_training_step: 1
15 general:
16   benchmark_csv_path: null
17   consumed_train_samples: null
18   ignore_sanity_checks: true
19   project: template
20   run: template
21   seed: 42
22   step: null
23 lighteval: null
24 logging:
25   iteration_step_info_interval: 1
26   log_level: info
27   log_level_replica: info
28 model:
29   ddp_bucket_cap_mb: 25
30   dtype: bfloat16
31   init_method:
32     std: 0.025
33   make_vocab_size_divisible_by: 1
34   model_config:
35     bos_token_id: 1
36     eos_token_id: 2
37     hidden_act: silu
38     hidden_size: 1536
39     initializer_range: 0.02
40     intermediate_size: 6144
41     is_llama_config: true
42     max_position_embeddings: 1024
43     num_hidden_layers: 24
44     num_attention_heads: 16
45     num_key_value_heads: 16
46     pad_token_id: null
47     pretraining_tp: 1

```

```

48     rms_norm_eps: 1.0e-06
49     rope_scaling: null
50     tie_word_embeddings: true
51     use_cache: true
52     vocab_size: 131072
53 optimizer:
54   optimizer_factory:
55     adam_beta1: 0.9
56     adam_beta2: 0.95
57     adam_eps: 1.0e-08
58     name: adamW
59     torch_adam_is_fused: true
60   learning_rate_scheduler:
61     learning_rate: 0.0008
62     lr_decay_starting_step: 61001 # for 119B tokens (36001 for 70B
        tokens, 15001 for 30B tokens)
63     lr_decay_steps: 12000 # for 119B tokens (7000 for 70B tokens, 4000
        for 30B tokens)
64     lr_decay_style: 1-sqrt
65     lr_warmup_steps: 2000
66     lr_warmup_style: linear
67     min_decay_lr: 0.00
68   zero_stage: 0
69   clip_grad: 1.0
70   weight_decay: 0.1
71   accumulate_grad_in_fp32: true
72 parallelism:
73   dp: 80
74   expert_parallel_size: 1
75   pp: 1
76   pp_engine: 1flb
77   tp: 1
78   tp_linear_async_communication: true
79   tp_mode: REDUCE_SCATTER
80 profiler: null
81 tokenizer:
82   tokenizer_max_length: null
83   tokenizer_name_or_path: mistralai/Mistral-Nemo-Base-2407
84   tokenizer_revision: null
85 tokens:
86   batch_accumulation_per_replica: 1
87   limit_test_batches: 0
88   limit_val_batches: 0
89   micro_batch_size: 20
90   sequence_length: 1024
91   train_steps: 73000 # for 119B tokens (43000 for 70B tokens, 19000 for
        30B tokens)
92   val_check_interval: -1

```

B ADDITIONAL RESULTS

B.1 MODEL SELECTION - PER LANGUAGE RESULTS

For completeness, we present the individual benchmark results of the 1B-parameter model trained on 119B tokens for each language in the following tables: Table 9 for Chinese, Table 10 for French, Table 11 for German, Table 12 for Arabic, and Table 13 for Danish.

Table 9: Benchmark performance comparison in Chinese between the baseline (FineWeb-2) and our proposed filtering methods (*FT*, *MLP*, and *CS*) trained on *MKC*⁺ or *MKC*, retaining 10% of the documents. The average rank is computed across FineTasks performance of 1B-parameter models evaluated after 119B tokens were consumed.

Approach	<i>MLP MKC</i> ⁺	<i>MLP MKC</i>	<i>CS MKC</i>	<i>FT MKC</i>	<i>FT MKC</i> ⁺	Baseline	<i>CS MKC</i> ⁺
Average Rank	1.7333	2.4333	4.0667	4.0667	4.4667	5.2333	6.0000
AGIEval	0.2995	0.2948	0.2897	0.2919	0.2817	0.2853	0.2773
Belebele	0.3300	0.3233	0.3178	0.3133	0.3133	0.3056	0.3022
C ³	0.4550	0.4480	0.4400	0.4500	0.4400	0.4400	0.4370
C-Eval	0.3095	0.3060	0.2760	0.2903	0.2906	0.2878	0.2805
CMMLU	0.3312	0.3259	0.3041	0.3043	0.3060	0.3009	0.2995
CMRC 2018	0.2224	0.2125	0.1614	0.2251	0.2164	0.1949	0.1866
HellaSwag	0.3790	0.3800	0.3530	0.3680	0.3660	0.3510	0.3370
M3Exam	0.3319	0.3245	0.3084	0.3201	0.3245	0.3216	0.3245
X-CODAH	0.3033	0.3000	0.3233	0.3100	0.2900	0.2967	0.3067
X-CSQA	0.2740	0.2680	0.2690	0.2610	0.2520	0.2510	0.2650
XCOPA	0.6200	0.6400	0.6180	0.5740	0.5740	0.6000	0.5620
OCNLI	0.5470	0.5470	0.5340	0.5250	0.5600	0.5420	0.5060
Chinese-SQuAD	0.0929	0.1097	0.0865	0.0889	0.0850	0.0777	0.0585
XStoryCloze	0.5800	0.5630	0.5710	0.5560	0.5610	0.5580	0.5570
XWINO	0.6429	0.6528	0.6587	0.6131	0.5992	0.6429	0.6111

Table 10: Benchmark performance comparison in French between the baseline (FineWeb-2) and our proposed filtering methods (*FT*, *MLP*, and *CS*) trained on *MKC*⁺ or *MKC*, retaining 10% of the documents. The average rank is computed across FineTasks performance of 1B-parameter models evaluated after 119B tokens were consumed.

Approach	<i>FT MKC</i> ⁺	<i>MLP MKC</i> ⁺	<i>MLP MKC</i>	<i>FT MKC</i>	<i>CS MKC</i>	<i>CS MKC</i> ⁺	Baseline
Average Rank	3.2222	3.5000	3.5556	3.7778	4.0000	4.6667	5.2778
Belebele	0.3378	0.3533	0.3678	0.3489	0.3444	0.3344	0.3444
HellaSwag	0.5380	0.5380	0.4990	0.5150	0.5280	0.5070	0.5180
X-CSQA	0.2820	0.2740	0.2730	0.2990	0.2850	0.2900	0.2870
XNLI 2.0	0.7340	0.7400	0.7430	0.7230	0.7450	0.7330	0.7180
FQuAD	0.2597	0.2803	0.3032	0.2981	0.2411	0.2476	0.2401
MMLU	0.2896	0.2895	0.2925	0.2886	0.2806	0.2815	0.2706
Mintaka	0.0710	0.0438	0.0334	0.0670	0.0610	0.0976	0.0712
X-CODAH	0.3000	0.2667	0.2867	0.2767	0.3000	0.2800	0.2633
ARC (Challenge)	0.3120	0.3180	0.3090	0.3060	0.2950	0.2830	0.2850

Table 11: Benchmark performance comparison in German between the baseline (FineWeb-2) and our proposed filtering methods (*FT*, *MLP*, and *CS*) trained on *MKC*⁺ or *MKC*, retaining 10% of the documents. The average rank is computed across FineTasks performance of 1B-parameter models evaluated after 119B tokens were consumed.

Approach	<i>MLP MKC</i> ⁺	<i>FT MKC</i> ⁺	<i>FT MKC</i>	<i>CS MKC</i>	<i>MLP MKC</i>	<i>CS MKC</i> ⁺	Baseline
Average Rank	3.1250	3.1250	3.5000	3.7500	4.5000	4.7500	5.2500
MMLU	0.2940	0.2879	0.2926	0.2770	0.2905	0.2764	0.2718
ARC (Challenge)	0.2760	0.2850	0.2820	0.2880	0.2830	0.2640	0.2680
Mintaka	0.0580	0.0548	0.0735	0.0576	0.0494	0.0766	0.0498
Belebele	0.3611	0.3578	0.3544	0.3544	0.3567	0.3422	0.3544
X-CODAH	0.3367	0.3500	0.3300	0.3567	0.3400	0.3600	0.3467
X-CSQA	0.2978	0.3008	0.2877	0.2887	0.2857	0.2918	0.2787
HellaSwag	0.4640	0.4710	0.4870	0.4820	0.4540	0.4390	0.4470
XNLI 2.0	0.6620	0.6530	0.6740	0.6440	0.6610	0.6520	0.6890

Table 12: Benchmark performance comparison in Arabic between the baseline (FineWeb-2) and our proposed filtering methods (*FT*, *MLP*, and *CS*) trained on *MKC*⁺ or *MKC*, retaining 56% of the documents. The average rank is computed across FineTasks performance of 1B-parameter models evaluated after 119B tokens were consumed.

Approach	<i>MLP MKC</i> ⁺	<i>MLP MKC</i>	<i>FT MKC</i> ⁺	Baseline	<i>CS MKC</i> ⁺	<i>CS MKC</i>	<i>FT MKC</i>
Average Rank	2.7812	3.2500	3.6875	3.9688	3.9688	5.0312	5.3125
EXAMS	0.3537	0.3656	0.3552	0.3582	0.3443	0.3262	0.3346
MMLU	0.4007	0.3909	0.4023	0.3894	0.3912	0.3781	0.3885
ARC (Easy)	0.4330	0.4230	0.4210	0.4120	0.4020	0.3940	0.4080
AlGhafa SciQ	0.6915	0.7005	0.6965	0.6854	0.6724	0.6683	0.6804
Belebele	0.3456	0.3356	0.3322	0.3311	0.3356	0.3567	0.3233
SOQAL	0.7333	0.6867	0.7000	0.7200	0.7267	0.6867	0.7133
MLQA	0.2386	0.2402	0.1928	0.1901	0.2189	0.2154	0.1793
TyDi QA	0.1547	0.1476	0.1230	0.1441	0.1223	0.1097	0.1182
AlGhafa RACE	0.3720	0.3740	0.3640	0.3710	0.3590	0.3660	0.3730
ARCD	0.3638	0.3505	0.3235	0.3354	0.3358	0.3432	0.3043
X-CODAH	0.2600	0.2533	0.2567	0.2633	0.2633	0.2500	0.2600
AlGhafa PIQA	0.6360	0.6320	0.6400	0.6240	0.6320	0.6320	0.6370
X-CSQA	0.2740	0.2810	0.2770	0.2900	0.2880	0.2720	0.2770
XNLI 2.0	0.6570	0.6910	0.6990	0.7010	0.6910	0.6900	0.6770
HellaSwag	0.4270	0.4220	0.4280	0.4250	0.4260	0.4320	0.4150
XStoryCloze	0.6150	0.6100	0.6100	0.6070	0.6130	0.6180	0.5930

Table 13: Benchmark performance comparison in Danish between the baseline (FineWeb-2) and our proposed filtering methods (*FT*, *MLP*, and *CS*) trained on *MKC*⁺ or *MKC*, retaining 65% of the documents. The average rank is computed across FineTasks performance of 1B-parameter models evaluated after 119B tokens were consumed.

Approach	<i>CS MKC</i> ⁺	<i>MLP MKC</i> ⁺	<i>FT MKC</i> ⁺	Baseline
Average Rank	1.0000	2.5000	3.1667	3.3333
ARC (Challenge)	0.2820	0.2650	0.2730	0.2560
HellaSwag	0.4950	0.4850	0.4750	0.4750
Belebele	0.3333	0.3289	0.3189	0.3289

B.2 THRESHOLD SELECTION

To confirm that the *CS* filtering method is not competitive with *MLP* and *FT*, even when a higher percentage of documents is retained, we present the complete threshold selection results, including the *CS* method, in Table 14 in addition to the results shown in Section 4.2.2 (Table 2).

We provide further results on the variation in the average length of documents retained by our model-based filtering approaches for Chinese, French, Arabic, and Danish. These results complement the findings for German discussed in Section 4.2.2 and are shown in Figure 3. Table 15 lists the actual dataset sizes (number of retained tokens) after tokenization for all languages.

B.3 TRAINING DATA ANALYSIS

We give details on the variation in the average length of documents retained by our model-based filtering method *MLP* for Chinese, French, Arabic, and Danish with different training datasets. The results are shown for German in Figure 4 and for all other languages in Figure 5.

B.4 IMPACT ON MULTILINGUAL MODEL TRAINING

This section presents the results of our *MLP MKC*⁺ approach on multilingual model training for Chinese (Table 16), Arabic (Table 17), German (Table 18), and Danish (Table 19), in addition to the results for French discussed in Section 4.2.7.

Table 14: Benchmark performance comparison (average rank) between the baseline (FineWeb-2) and our proposed filtering methods (*FT*, *MLP*, *CS*) trained on MKC^+ or MKC , retaining top 10%, 15% or 20% of the documents. The average rank is computed across FineTasks performance of 1B-parameter models evaluated for Chinese, German and French after 70B and 119B tokens were consumed.

Approach	Threshold	Average Rank
<i>MLP MKC</i> ⁺	10%	11.73
<i>MLP MKC</i> ⁺	15%	12.13
<i>MLP MKC</i>	20%	15.07
<i>MLP MKC</i>	15%	15.09
<i>MLP MKC</i> ⁺	20%	15.40
<i>MLP MKC</i>	10%	16.09
<i>FT MKC</i> ⁺	10%	18.61
<i>CS MKC</i>	15%	19.02
<i>CS MKC</i>	20%	19.24
<i>FT MKC</i>	15%	19.84
<i>FT MKC</i>	10%	20.02
<i>CS MKC</i>	10%	20.67
<i>FT MKC</i>	20%	20.80
<i>FT MKC</i> ⁺	15%	22.05
<i>FT MKC</i> ⁺	20%	22.52
<i>CS MKC</i> ⁺	15%	24.66
<i>CS MKC</i> ⁺	20%	25.08
Baseline	—	25.54
<i>CS MKC</i> ⁺	10%	26.94




Figure 3: Comparison of average document length and standard deviation in FineWeb-2 before and after filtering using one of our approaches retaining top 10% of the documents for Chinese and French, 56% for Arabic and 65% for Danish. The average document length of FineWeb-2 is represented as a red horizontal line, while the medians are shown as red dots. Document length is measured based on number of space-separated tokens.

Table 15: Comparison of retained tokens in FineWeb-2 before and after filtering using one of our proposed approaches retaining top 10% of the documents for Chinese, French and German, 56% for Arabic and 65% for Danish. The token counts correspond to the size of the tokenized datasets, processed with the multilingual Mistral v3 (Tekken) tokenizer (Mistral AI, 2024).

Approach	Chinese	French	German	Arabic	Danish
<i>MLP MKC</i> ⁺	150B (9%)	89B (12%)	119B (12%)	78B (61%)	71B (66%)
<i>MLP MKC</i>	105B (7%)	72B (10%)	87B (9%)	75B (59%)	–
<i>FT MKC</i> ⁺	221B (14%)	70B (10%)	63B (6%)	77B (61%)	70B (65%)
<i>FT MKC</i>	190B (12%)	43B (6%)	65B (7%)	80B (63%)	–
<i>CS MKC</i> ⁺	170B (11%)	126B (17%)	166B (17%)	82B (65%)	77B (71%)
<i>CS MKC</i>	161B (10%)	132B (18%)	172B (18%)	83B (65%)	–
Baseline	1597B	730B	973B	127B	108B




Figure 4: Comparison of average document length and standard deviation in FineWeb-2 before and after filtering using *MLP* filtering method retaining top 10% of the documents with different training datasets. The average document length of FineWeb-2 is represented as a red horizontal line, while the medians are shown as red dots. Document length is measured based on number of space-separated tokens.




Figure 5: Comparison of average document length and standard deviation in FineWeb-2 before and after filtering using *MLP* filtering method retaining top 10% of the documents for Chinese and French, 56% for Arabic and 65% for Danish with different training datasets. The average document length of FineWeb-2 is represented as a red horizontal line, while the medians are shown as red dots. Document length is measured based on number of space-separated tokens.

Table 16: Benchmark performance comparison for Chinese of multilingual LLMs (M) trained on FineWeb-2 or the refined dataset using our $MLP MKC^+$ approach (retaining top 10% of the documents for Chinese, German, and French, 56% for Arabic, and 65% for Danish) trained on 595B tokens, against their monolingual counterparts trained on 119B tokens. The average rank is computed across FineTasks performance for 1B-parameter models trained on 119B tokens.

Dataset	Ours	Ours $_M$	FW-2 $_M$	FW-2
Average Rank	1.5667	2.1667	2.9000	3.3667
AGIEval	0.2995	0.2863	0.2894	0.2853
Belebele	0.3300	0.3456	0.3189	0.3056
C ³	0.4550	0.4520	0.4480	0.4400
C-Eval	0.3095	0.2848	0.2683	0.2878
CMMLU	0.3312	0.3064	0.2967	0.3009
CMRC 2018	0.2224	0.2689	0.2090	0.1949
HellaSwag	0.3790	0.3740	0.3740	0.3510
M3Exam	0.3319	0.3040	0.3304	0.3216
X-CODAH	0.3033	0.3067	0.2800	0.2967
X-CSQA	0.2740	0.2810	0.2780	0.2510
XCOPA	0.6200	0.6020	0.5860	0.6000
OCNLI	0.5470	0.5320	0.4910	0.5420
Chinese-SQuAD	0.0929	0.1304	0.1017	0.0777
XStoryCloze	0.5800	0.5760	0.5650	0.5580
XWINO	0.6429	0.6409	0.6468	0.6429

Table 17: Benchmark performance comparison for Arabic of multilingual LLMs (M) trained on FineWeb-2 or the refined dataset using our $MLP MKC^+$ approach (retaining top 10% of the documents for Chinese, German, and French, 56% for Arabic, and 65% for Danish) trained on 595B tokens, against their monolingual counterparts trained on 119B tokens. The average rank is computed across FineTasks performance for 1B-parameter models trained on 119B tokens.

Dataset	Ours $_M$	Ours	FW-2	FW-2 $_M$
Average Rank	1.9688	2.0000	2.7500	3.2812
EXAMS	0.3336	0.3537	0.3582	0.3076
MMLU	0.3828	0.4007	0.3894	0.3599
ARC (Easy)	0.4190	0.4330	0.4120	0.3760
AlGhafa SciQ	0.6764	0.6915	0.6854	0.6563
Belebele	0.3511	0.3456	0.3311	0.3344
SOQAL	0.7000	0.7333	0.7200	0.6533
MLQA	0.2208	0.2386	0.1901	0.2085
TyDi QA	0.1634	0.1547	0.1441	0.1429
AlGhafa RACE	0.3830	0.3720	0.3710	0.3770
ARCD	0.3377	0.3638	0.3354	0.2970
X-CODAH	0.2767	0.2600	0.2633	0.2767
AlGhafa PIQA	0.6170	0.6360	0.6240	0.6160
X-CSQA	0.2860	0.2740	0.2900	0.2660
XNLI 2.0	0.7080	0.6570	0.7010	0.7340
HellaSwag	0.4390	0.4270	0.4250	0.4240
XStoryCloze	0.6370	0.6150	0.6070	0.6160

Table 18: Benchmark performance comparison for German of multilingual LLMs (M) trained on FineWeb-2 or the refined dataset using our $MLP\ MKC^+$ approach (retaining top 10% of the documents for Chinese, German, and French, 56% for Arabic, and 65% for Danish) trained on 595B tokens, against their monolingual counterparts trained on 119B tokens. The average rank is computed across FineTasks performance for 1B-parameter models trained on 119B tokens.

Dataset	Ours $_M$	Ours	FW-2	FW-2 $_M$
Average Rank	1.5000	2.1250	2.9375	3.4375
MMLU	0.2918	0.2940	0.2718	0.2691
ARC (Challenge)	0.2740	0.2760	0.2680	0.2640
Mintaka	0.0821	0.0580	0.0498	0.0660
Belebele	0.3956	0.3611	0.3544	0.3633
X-CODAH	0.3500	0.3367	0.3467	0.3167
X-CSQA	0.3048	0.2978	0.2787	0.2787
HellaSwag	0.4690	0.4640	0.4470	0.4430
XNLI 2.0	0.6420	0.6620	0.6890	0.6340

Table 19: Benchmark performance comparison for Danish of multilingual LLMs (M) trained on FineWeb-2 or the refined dataset using our $MLP\ MKC^+$ approach (retaining top 10% of the documents for Chinese, German, and French, 56% for Arabic, and 65% for Danish) trained on 595B tokens, against their monolingual counterparts trained on 119B tokens. The average rank is computed across FineTasks performance for 1B-parameter models trained on 119B tokens.

Dataset	Ours $_M$	Ours	FW-2 $_M$	FW-2
Average Rank	1.6667	2.1667	3.0000	3.1667
ARC (Challenge)	0.2920	0.2650	0.2600	0.2560
HellaSwag	0.4710	0.4850	0.4560	0.4750
Belebele	0.3700	0.3289	0.3311	0.3289

C LIST OF EVALUATION BENCHMARKS AND METRICS

We provide a detailed overview of the evaluation benchmarks used to assess our models’ performance, along with their respective evaluation metrics in Table 20. For non-English tasks and English MMLU, we use the *cloze* multiple-choice prompt, which allows the model to directly predict each option instead of using the standard prompt format with A/B/C/D letter prefixes as targets. This approach was chosen because it has been shown to serve as a more reliable performance indicator earlier in training (Kydliček et al., 2024). We evaluate the models in a 0-shot setting.

Table 20: List of evaluation benchmarks and metrics used in our setup for Chinese, French, German, Arabic, Danish, and English.

Benchmark	Chinese	French	German	Arabic	Danish	English	Evaluation metric
AGIEval (Zhong et al., 2023)	✓						Normalized accuracy
AlGhafa ARC (Almazrouei et al., 2023)				✓			Normalized accuracy
AlGhafa PIQA (Almazrouei et al., 2023)				✓			Normalized accuracy
AlGhafa RACE (Almazrouei et al., 2023)				✓			Normalized accuracy
AlGhafa SciQ (Almazrouei et al., 2023)				✓			Normalized accuracy
ArabicMMLU (Koto et al., 2024)				✓			Normalized accuracy
ARC (Clark et al., 2018)						✓	Normalized accuracy
ARCD (Mozannar et al., 2019)				✓			F1 score
Belebele (Bandarkar et al., 2024)	✓	✓	✓	✓	✓		Normalized accuracy
C ³ (Sun et al., 2020)	✓						Normalized accuracy
C-Eval (Huang et al., 2023)	✓						Normalized accuracy
Chinese-SQuAD (Pluto-Junzeng, 2019)	✓						F1 score
CMMLU (Li et al., 2024a)	✓						Normalized accuracy
CMRC 2018 (Cui et al., 2019)	✓						F1 score
CommonsenseQA (Talmor et al., 2019)						✓	Normalized accuracy
EXAMS (Hardalov et al., 2020)				✓			Normalized accuracy
FQuAD (d’Hoffschmidt et al., 2020)		✓					F1 score
HellaSwag (Zellers et al., 2019)						✓	Normalized accuracy
M3Exam (Zhang et al., 2023)	✓						Normalized accuracy
Meta MMLU (Llama Team, 2024)		✓	✓				Normalized accuracy
Mintaka (Sen et al., 2022)		✓	✓				F1 score
MLMM ARC (Lai et al., 2023)		✓	✓		✓		Normalized accuracy
MLMM HellaSwag (Lai et al., 2023)	✓	✓	✓	✓	✓		Normalized accuracy
MLQA (Lewis et al., 2020)				✓			F1 score
MMLU (Hendrycks et al., 2020)						✓	Normalized accuracy
OCNLI (Hu et al., 2020)	✓						Normalized accuracy
OpenBookQA (Mihaylov et al., 2018)						✓	Normalized accuracy
PIQA (Bisk et al., 2019)						✓	Normalized accuracy
SOQAL (Mozannar et al., 2019)				✓			Normalized accuracy
TriviaQA (Joshi et al., 2017)						✓	Quasi-exact match
TyDi QA (Clark et al., 2020)				✓			F1 score
WinoGrande (Sakaguchi et al., 2019)						✓	Normalized accuracy
X-CODAH (Lin et al., 2021a)	✓	✓	✓	✓			Normalized accuracy
XCOPA (Ponti et al., 2020)	✓						Normalized accuracy
X-CSQA (Lin et al., 2021a)	✓	✓	✓	✓			Normalized accuracy
XNLI 2.0 (Upadhyay & Upadhyay, 2023)		✓	✓	✓			Normalized accuracy
XStoryCloze (Lin et al., 2021b)	✓			✓			Normalized accuracy
XWINO (Tikhonov & Ryabinin, 2021)	✓						Normalized accuracy

D FINEWEB DOCUMENTS IN DIFFERENT SCORING APPROACHES

To illustrate the types of documents each classifier scores highly or poorly, we present the highest- and lowest-scoring FineWeb examples for each of our classifier approaches (*FT MKC⁺*, *MLP MKC⁺*, *CS MKC⁺*). These examples were selected from the randomly chosen FineWeb test dataset (10K) used to validate the training of our model-based classifiers.

D.1 FASTTEXT CLASSIFIER (FT)

Highest score:

hi. i couldn't solve my problem because it has two conditional logical propositions. the problem is:can anyone help me about this, thanks =)we're expected to know that: . is equivalent tofind a logically equivalent proposition for:by first writing its contrapositive, and then applying demorgan's lawand the equality forthey were trying to be helpful by outlining the steps we should follow,. . but i think they made it more confusing.i don't see the purpose of using the contrapositive here.. . i wouldn't have done it that way.besides, the statement is a tautology . . .which gives us: .and this is a tautology: "a thing implies itself" ... which is always true.i don't know of any "logically equivalent proposition" we can write . . .

Lowest score:

|starts||23 sep 2016 (fri) (one day only)|want to travel soon but dont wish to fork out a fortune for flights? check out todays promotion from jetstar featuring promo fares fr \$35 all-in valid for travel period commencing 12 october 2016dont miss out! all-in frenzy fares to hong kong, penang and more from \$35.sale ends 23 sep, 11pm! travelling||price||travel period||find flight||penang||\$35^|| [...]

D.2 MULTI-LAYER PERCEPTRON (MLP)

Highest score:

Naqadeh County is a county in West Azerbaijan Province in Iran. The capital of the county is Naqadeh. At the 2006 census, the county's population was 117,831, in 27,937 families. The county is subdivided into two districts: the Central District and Mohammadyar District. The county has two cities: Naqadeh and Mohammadyar.

Lowest score:

Custom Wedding Gifts
Personalized photo frames, albums & keepsakes. Heirloom quality!
Custom Engraved Journals
Handmade in Florence Italy. Dozens of sizes and paper styles!
Awesome Leather Journals
Personalized, Customizable, Artisan made in Santa Fe, NM.
Ink Rendering from Photos
100% Hand painted with unique style by pro artists. From \$49.

D.3 COSINE SIMILARITY (CS)

Highest score:

When you are renting a 5, 10, 15, 20, 30 or 40 yard dumpster, you want a company you can trust with prices that make you smile. Give us a call today and see the difference we can make in your next construction or clean out project.
Simply give us a call and we will help you figure out your dumpster rental needs.
Our dumpsters usually go out same-day or next-day depending on when you call.
We provide top-notch service, while going easy on your bottom line. What more could you ask for?
Our trained operators are here to give you a fast and hassle-free experience from start to finish.[...]

Lowest score:

Cooperative flat 206/J
- Cooperative flat 201/J - Sold
2(1)+kitchenette, 50,1 m2Cooperative flat 202/J - Sold
2(1)+kitchenette, 44,9 m2Cooperative flat 203/J - Sold
2(1)+kitchenette, 50,6 m2Cooperative flat 204/J - Sold
1+kitchenette, 27,1 m2Cooperative flat 205/J - Sold
2(1)+kitchenette, 50,1 m2Cooperative flat 206/J - On sale
3+kitchenette 86,7 m2[...]

E EXAMPLE OF A CONTAMINATED DOCUMENT

We present an example of a FineWeb document that was removed during our decontamination pipeline.

MMLU contaminated document (matched 13-gram in bold):

Here is our diagram of the Preamble to the Constitution of the United States. It is based on our understanding of the use of "in order to" as a subordinating conjunction that introduces a series of infinitival clauses (without subjects) that, in turn, modify the compound verbs "do ordain" and "establish."
See A Grammar of Contemporary English by Randolph Quirk, Sidney Greenbaum, Geoffrey Leech, and Jan Svartvik. Longman Group: London. 1978. p. 753.
We the People of the United States, in Order to form a more perfect Union, establish Justice, insure domestic Tranquility, **provide for the common defence, promote the general Welfare, and secure the Blessings** of Liberty to ourselves and our Posterity, do ordain and establish this Constitution for the United States of America.
If you have alternative rendering for this sentence, we would be happy to hear of it. Use the e-mail icon to the left.