

What Drives Performance in Multilingual Language Models?

Sina Bagheri Nezhad, Ameeta Agrawal
Portland State University
{sina.bagherinezhad, ameeta}@pdx.edu

Abstract

This study investigates the factors influencing the performance of multilingual large language models (MLLMs) across diverse languages. We study 6 MLLMs, including masked language models, autoregressive models, and instruction-tuned LLMs, on the SIB-200 dataset, a topic classification dataset encompassing 204 languages. Our analysis considers three scenarios: ALL languages, SEEN languages (present in the model’s pretraining data), and UNSEEN languages (not present or documented in the model’s pretraining data in any meaningful way). We examine the impact of factors such as pretraining data size, general resource availability, language family, and script type on model performance. Decision tree analysis reveals that pretraining data size is the most influential factor for SEEN languages. However, interestingly, script type and language family are crucial for UNSEEN languages, highlighting the importance of cross-lingual transfer learning. Notably, model size and architecture do not significantly alter the most important features identified. Our findings provide valuable insights into the strengths and limitations of current MLLMs and hope to guide the development of more effective and equitable multilingual NLP systems.¹

1 Introduction

Multilingual large language models (MLLMs) have revolutionized natural language processing by enabling applications like machine translation and sentiment analysis across numerous languages (Barbieri et al., 2022; Yang et al., 2023). Understanding how these models perform across languages with diverse linguistic properties is crucial for further development (Devlin et al., 2019; Wu and Dredze, 2020; Scao et al., 2022; Lai et al., 2023; Ahuja et al., 2023). Despite significant

progress, linguistic disparities persist in NLP, highlighting the need for models that perform effectively and safely across a wider range of languages (Joshi et al., 2020; Ranathunga and de Silva, 2022; Agrawal et al., 2023; Wang et al., 2023).

The factors contributing to the effectiveness of MLLMs, however, remain unclear. While several studies suggest the amount of language-specific pretraining data as a key factor (Wu and Dredze, 2020; Scao et al., 2022; Shliazhko et al., 2022; Ahuja et al., 2023), most investigations are limited in scope, focusing on a small set of languages, specific tasks, or training paradigms like masked language modeling (MLM) or autoregressive models. Crucially, prior work often overlooks the distinction between languages encountered during pretraining (SEEN), languages entirely new to the model (UNSEEN), and the complete set of languages available in the evaluation dataset (ALL). The question remains – *what factors are important in the case of unseen languages where language-specific pretraining data is not one of the relevant factors?* This distinction is essential for understanding how MLLMs generalize to languages with varying levels of familiarity.

Our work takes a deeper look at the various factors under several experimental settings. Our key contributions are as follows:

- We conduct a comprehensive evaluation of 6 MLLMs, including MLM, autoregressive, and instruction-tuned LLMs, on a text classification task spanning a wide range of languages. This diverse set of models includes mBERT (Devlin et al., 2019), XLM-R (Conneau et al., 2020), GPT-3.5 (Brown et al., 2020), Bloom (Scao et al., 2022) in 5 sizes, Bloomz (Muennighoff et al., 2023) in 5 sizes, and XGLM (Lin et al., 2022) in 4 sizes. Additionally, we consider three training scenarios: zero-shot, 2-shot, and fully supervised.
- We consider four key factors in our analysis: pre-

¹https://github.com/PortNLP/MLLMs_performance

Reference	Factors	Task	Languages
Wu and Dredze (2020)	Pretraining data size, Task-specific data size, Vocabulary size	NER	99
Scao et al. (2022)	Pretraining data size, Task-specific data size, Language family, Language script	Probing	17
Shliazhko et al. (2022)	Pretraining data size, Language script, Model size	Perplexity	61
Ahuja et al. (2023)	Pretraining data size, Tokenizer fertility	Classification, QA, Sequence Labeling, NLG	2-48
Ours	Pretraining data size, Language family, Language script, General resource availability	Text classification	204

Table 1: Factors considered in related works and this work.

training data size, general resource availability levels, language family, and script type. This allows for a more nuanced understanding of the factors influencing MLLM performance.

- We leverage the recently introduced SIB-200 dataset (Adelani et al., 2023), which includes 204 languages, enabling us to investigate MLLM performance across a diverse and extensive linguistic landscape. Between the languages pertaining to the models and the dataset, we are able to further distinguish them along the dimensions of SEEN, UNSEEN, or ALL, depending on whether the languages were seen during pretraining, or unseen during pretraining, or the set of all languages available in the evaluation dataset, respectively.

By analyzing these factors across different models and training setups, we aim to provide deeper insights into the development of effective and equitable MLLMs for a truly multilingual NLP landscape.

2 Related Work

Multilingual NLP research has flourished in recent years, with the development and evaluation of numerous multilingual language models trained on diverse and extensive language datasets. Notable examples include mBERT (Devlin et al., 2019), XLM-R (Conneau et al., 2020), mBART (Liu et al., 2020), mT5 (Xue et al., 2021), BLOOM (Scao et al., 2022), GPT-3 (Brown et al., 2020), GPT-4 (OpenAI, 2023), LLaMA (Touvron et al., 2023), PaLM (Chowdhery et al., 2022), and PaLM 2 (Anil et al., 2023).

Researchers are increasingly interested in investigating the factors influencing MLLM performance.

Wu and Dredze (2020) examined the impact of pretraining data size, task-specific data size, and vocabulary size on named entity recognition performance. Scao et al. (2022) explored the correlation between probing performance and factors like language family, task-specific dataset size, and pretraining dataset size for the BLOOM model. Shliazhko et al. (2022) assessed the impact of language script, pretraining corpus size, and model size on language modeling performance, while Ahuja et al. (2023) investigated the influence of tokenizer fertility and pretraining data on MLLM performance.

While these studies provide valuable insights, they often focus on a limited set of languages, primarily due to the historical scarcity of annotated multilingual datasets. Additionally, research by Blasi et al. (2022) highlights the significant inequalities in the development and performance of language technologies across the world’s languages, with a strong bias towards resource-rich languages like English and other Western European languages. Further exacerbating this issue is the lack of representation for dialects, varieties, and closely-related languages within existing datasets. As noted by Faisal et al. (2024), this absence hinders the development of NLP systems capable of effectively handling the nuances of linguistic diversity. However, the recent emergence of comprehensive multilingual datasets like SIB-200 (Adelani et al., 2023), and GLOT500 (ImaniGooghari et al., 2023) offers exciting opportunities for more extensive and nuanced analyses. Table 1 summarizes the factors considered in related works and our study. For a more comprehensive overview of contributing factors to cross-lingual transfer in multilingual language models, readers are encouraged to refer to the review by Philipppy et al. (2023).

3 Methodology

Several factors can influence the performance of multilingual models. In this section, we briefly describe the distinct factors related to typology and data, the dataset of more than 200 languages used for evaluation, and the models we consider in this study.

3.1 Typology and Data Factors

We consider various factors to understand their impact on model performance including:

- **Pretraining Data Size:** This refers to the percentage of language-specific data used during the pretraining of each model².
- **General Resource Availability (Res Level):** Beyond model-specific resources such as pretraining data size, we also consider a more general notion of resource availability, as per the linguistic diversity taxonomy which categorizes languages into six resource levels (Joshi et al., 2020), where level 0 corresponds to low-resource and level 5 corresponds to high-resource level languages. This classification helps us understand the influence of more general resource availability on model performance, and may serve as a proxy when model-specific statistics may not be available (such as in the case of proprietary models). Language resource levels generally correlate positively with models pretraining data sizes, with varying degrees of alignment across different models: mBERT (0.52) and XLM-R (0.48) exhibit relatively stronger correlations, while GPT-3 (0.18), BLOOM (0.37), and XGLM (0.31) show comparatively weaker associations.
- **Language Family (Lang Family):** The language families that the languages belong to capture some of their linguistic relationships. The information was sourced from the Ethnologue³ (Ethnologue, 2022).
- **Script:** The script of a language refers to the writing system it employs. This information was sourced from ScriptSource⁴.

²We obtained the train dataset distribution values for mBERT from <https://github.com/mayhewsw/multilingual-data-stats> and for GPT-3.5 we use proxy statistics from https://github.com/openai/gpt-3/blob/master/dataset_statistics/languages_by_word_count.csv. Distribution of train dataset for XLM-R, BLOOM, BLOOMZ and XGLM were obtained from their respective papers.

³<https://www.ethnologue.com>

⁴<https://www.scriptsource.org>

3.2 Data

We systematically study the multilingual models under an important NLP task – text classification (Chang and Bergen, 2023). The SIB-200 dataset (Adelani et al., 2023) offers a valuable resource for evaluating MLLM performance in a large-scale text classification task, enabling simultaneous analysis of approximately 200 languages, with text samples categorized into one of seven classes. F1 score is used as the metric for this task.

Exploratory analysis of the dataset reveals several interesting insights:

- As shown in Figure 1, most languages in SIB-200 are classified as resource level 1, indicating a deliberate focus on low-resource languages. This allows us to assess how MLLMs perform on languages with limited linguistic resources available.
- Figure 4 in Appendix B illustrates the distribution of language families within the SIB-200 dataset. Notably, the dataset encompasses 23 different language families, providing a rich linguistic landscape for our analysis. Indo-European languages constitute a significant portion (approximately 36%) of SIB-200, reflecting their status as the most widely spoken language family globally (Ethnologue, 2022). However, Niger-Congo, Afro-Asiatic, and Austronesian languages also have considerable representation in the dataset. This diverse language family distribution enables us to analyze MLLM performance across different linguistic groups.
- The SIB-200 dataset encompasses text samples written in 29 different script types, offering a diverse range of writing systems for our analysis. As shown in Figure 5 in Appendix B, the Latin script, used by nearly 70% of the global population (Vaughan, 2020), is the most prevalent writing system in the dataset, followed by Arabic and Cyrillic scripts. This distribution allows us to investigate the impact of script type on MLLM performance.

For all evaluations, we use the default train and test splits recommended by the SIB-200 authors. This ensures consistency and comparability across different models and training settings.

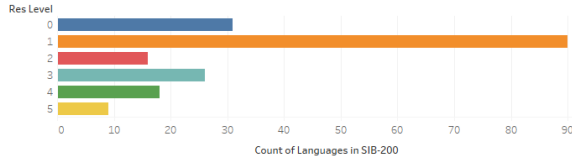


Figure 1: Distribution of resource levels in SIB-200.

3.3 Models

We study the following 6 multilingual language models spanning various architectures and sizes:

- Masked Language Models (MLMs):
 - mBERT (bert-base-multilingual-cased) (Devlin et al., 2019)
 - XLM-R (xlm-roberta-base) (Conneau et al., 2020)
- Autoregressive Language Models
 - GPT-3.5 (text-davinci-003) (Brown et al., 2020)
 - Bloom (Scao et al., 2022) in 5 sizes (560m, 1.1b, 1.7b, 3b, and 7.1b parameters)
 - XGLM (Lin et al., 2022) in 4 sizes (564m, 1.7b, 2.9b, and 7.5b parameters)
- Instruction-tuned LLMs:
 - Bloomz (Muennighoff et al., 2023) in 5 sizes (560m, 1.1b, 1.7b, 3b, and 7.1b parameters)

These models were chosen for several key reasons:

1. These models provide broad language coverage, allowing us to analyze performance across a diverse set of languages and maximize the linguistic diversity in our study.
2. By including MLMs, autoregressive models, and instruction-tuned LLMs, we can investigate how different model architectures influence performance.
3. The inclusion of models with varying parameter sizes allows us to investigate the interplay between model scale and the factors influencing performance.
4. mBERT and XLM-R, despite being relatively smaller models, have demonstrated competitive performance even compared to larger models like ChatGPT after fine-tuning (Lai et al., 2023; Zhu et al., 2023).

5. The inclusion of both Bloom and XGLM, both autoregressive models, allows us to investigate the impact of pretraining data composition. Bloom focuses more on low-resource languages during pretraining, whereas XGLM emphasizes high-resource languages. This deliberate selection enables us to analyze how the distribution of languages in the pretraining data affects performance across different resource levels.

Note that we primarily focus on models that are open-source or have made the list of pretraining languages and data composition available.

Additionally, we consider the following training and inference scenarios:

- Zero-shot: GPT-3.5, Bloom, Bloomz, and XGLM were evaluated directly on the test set without any specific fine-tuning. This assesses the model’s ability to generalize to unseen tasks and languages based on its pretrained knowledge.
- Two-shot In-Context Learning (ICL): Bloom, Bloomz, and XGLM were also evaluated in two-shot ICL setting where the models were provided with two labeled examples for each class from the train set. This allows us to particularly investigate effective factors for improving performance of unseen languages. We opted for two demonstrations in ICL to keep the input length shorter than the context length of our models across all languages.
- Full-shot: mBERT and XLM-R were fine-tuned on the SIB-200 training set and evaluated on the test set.

For full-shot training of mBERT and XLM-R, we adhered to the hyperparameters recommended by the SIB-200 paper authors to ensure consistency with the original dataset benchmarks. For Bloom, Bloomz, and XGLM in both zero-shot and two-shot ICL settings, as well as for GPT-3.5 in zero-shot setting, we use prompts to frame the text classification task, which are detailed in Appendix A.

4 Results and Analysis

Now we discuss the results of our comprehensive experiments. We focus on analyzing the performance of models across three distinct scenarios: ALL, SEEN, and UNSEEN. The ALL

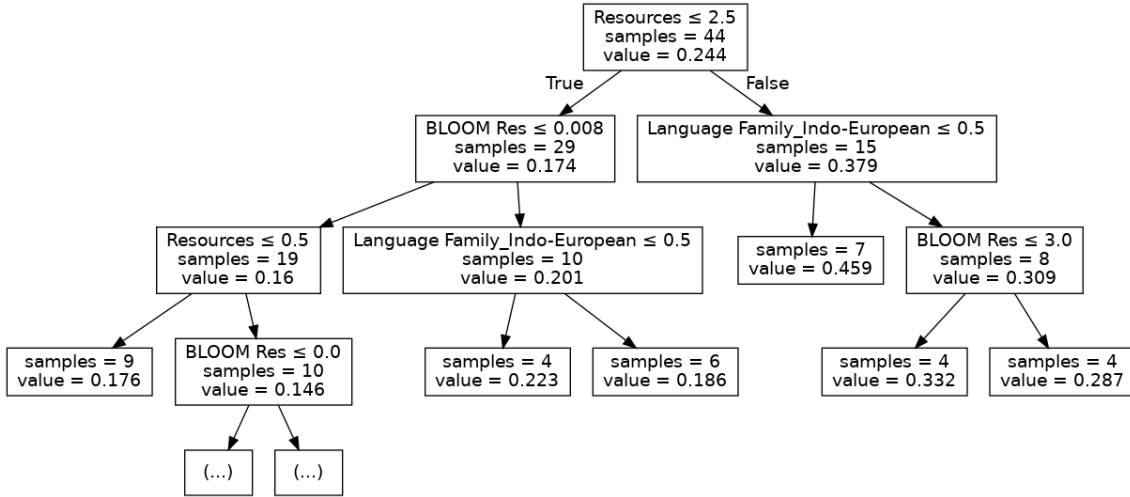


Figure 2: Decision tree for Bloom-560m (zero-shot, SEEN languages). “General resource level“ emerges as the most important feature, with a significant performance difference between languages above and below the 2.5 threshold ($p < 0.001$ as per Mann-Whitney U test).

scenario considers all languages in the SIB-200 dataset for which resource level information is available⁵. The SEEN scenario focuses on languages included in the pretraining data of the respective MLLMs, while the UNSEEN scenario examines performance on languages not present in the pretraining data.

In total, results are obtained from 93 distinct experimental settings (models of different sizes, training scenarios, and language categories of seen/unseen/all).

To understand the complex interplay of multiple factors influencing MLLM performance, we employ decision tree analysis for statistical inference. This approach is well-suited for handling factors of different types, including categorical, ordinal, and numeric data. Decision trees are trained to predict the F1 score of models based on language features. By analyzing the resulting tree structure, we can gain insights into the relative importance of different features and their interactions.

As decision trees were trained on the entirety of our data, traditional methods for testing their performance were not applicable. Instead, we employed the Mann-Whitney U test (Mann and Whitney, 1947), to ensure that the features appearing at the root of the decision trees were indeed relevant and contributed significantly to the differentiation between the language splits. This approach allowed us to validate the significance of the features identified by the decision tree in delineating

distinct language groups without relying solely on the performance metrics of the decision tree models themselves.

Figure 2 presents the decision tree analysis for the Bloom-560m model on SEEN languages, revealing *general resource level* as the most influential feature. Specifically, the tree distinguishes between languages with resource levels below 2.5 (levels 0,1,2) and those above 2.5 (levels 3,4,5). Among the 44 SEEN languages, the 29 languages with resource levels below 2.5 exhibit a mean F1 score of 0.174, while the 15 languages with higher resource levels achieve a significantly higher mean F1 score of 0.379. A Mann-Whitney U test confirms a statistically significant difference in performance between these two groups ($p < 0.001$). This suggests that for the Bloom-560m model on SEEN languages, the general resource level of a language plays a crucial role in determining its performance, with higher resource levels leading to better performance. By employing this combined approach of decision tree analysis and statistical testing, we can effectively disentangle the complex relationships between various factors and their impact on MLLM performance.

The summarized results⁶ of all 93 decision tree analyses are presented in Table 2. We observe distinct patterns in feature importance across the three scenarios:

⁵This information is available for 190 languages.

⁶Detailed decision trees for all models and setups are available in our repository: https://github.com/PortNLP/MLLMs_performance

Zero-shot			
Model	ALL	SEEN	UNSEEN
Bloom-560m	Pretrain data ($\leq 0.125\%$)	Resource level (≤ 2.5)	Script (Latin or not)
Bloom-1b1	Pretrain data ($\leq 0.125\%$)	Resource level (≤ 2.5)	Script (Devanagari or not)
Bloom-1b7	Pretrain data ($\leq 0.175\%$)	Resource level (≤ 2.5)	Script (Latin or not)
Bloom-3b	Pretrain data ($\leq 0.175\%$)	Resource level (≤ 2.5)	Script (Latin or not)
Bloom-7b1	Pretrain data ($\leq 0.125\%$)	Resource level (≤ 2.5)	Script (Devanagari or not)
Bloomz-560m	Script (Latin or not)	Pretrain data ($\leq 0.03\%$)	Script (Latin or not)
Bloomz-1b1	Pretrain data ($\leq 0.008\%$)	Pretrain data ($\leq 0.03\%$)	Script (Latin or not)
Bloomz-1b7	Pretrain data ($\leq 0.008\%$)	Pretrain data ($\leq 0.03\%$)	Script (Latin or not)
Bloomz-3b	Pretrain data ($\leq 0.002\%$)	Pretrain data ($\leq 0.013\%$)	Script (Latin or not)
Bloomz-7b1	Pretrain data ($\leq 0\%$)	Pretrain data ($\leq 0.9\%$)	Script (Latin or not)
XGLM-564m	Pretrain data ($\leq 0.003\%$)	Resource level (≤ 2)	Lang. family (Austronesian or not)
XGLM-1.7b	Pretrain data ($\leq 0.006\%$)	Pretrain data ($\leq 1.487\%$)	Script (Devanagari or not)
XGLM-2.9b	Pretrain data ($\leq 0.003\%$)	Script (Latin or not)	Script (Devanagari or not)
XGLM-7.5b	Pretrain data ($\leq 0\%$)	Pretrain data ($\leq 1.122\%$)	Script (Devanagari or not)
GPT-3.5	Resource level (≤ 2.5)	Pretrain data ($\leq 0.003\%$)	Lang. family (Indo-European or not)
Two-shot ICL			
Model	ALL	SEEN	UNSEEN
Bloom-560m	Pretrain data ($\leq 0.045\%$)	Pretrain data ($\leq 0.045\%$)	Lang. family (Indo-European or not)
Bloom-1b1	Pretrain data ($\leq 0.095\%$)	Pretrain data ($\leq 0.095\%$)	Script (Latin or not)
Bloom-1b7	Pretrain data ($\leq 0.175\%$)	Pretrain data ($\leq 0.175\%$)	Script (Latin or not)
Bloom-3b	Pretrain data ($\leq 0.008\%$)	Pretrain data ($\leq 0.008\%$)	Script (Latin or not)
Bloom-7b1	Pretrain data ($\leq 0.008\%$)	Pretrain data ($\leq 0.008\%$)	Script (Latin or not)
Bloomz-560m	Pretrain data ($\leq 0.03\%$)	Pretrain data ($\leq 0.03\%$)	Script (Devanagari or not)
Bloomz-1b1	Pretrain data ($\leq 0.008\%$)	Pretrain data ($\leq 0.013\%$)	Script (Latin or not)
Bloomz-1b7	Pretrain data ($\leq 0.005\%$)	Pretrain data ($\leq 0.013\%$)	Script (Cyrillic or not)
Bloomz-3b	Pretrain data ($\leq 0\%$)	Pretrain data ($\leq 0.9\%$)	Script (Latin or not)
Bloomz-7b1	Pretrain data ($\leq 0\%$)	Pretrain data ($\leq 0.013\%$)	Script (Latin or not)
XGLM-564m	Pretrain data ($\leq 0.003\%$)	Pretrain data ($\leq 0.095\%$)	Lang. family (Niger-Congo or not)
XGLM-1.7b	Pretrain data ($\leq 0.003\%$)	Resource level (≤ 2)	Script (Devanagari or not)
XGLM-2.9b	Pretrain data ($\leq 0.003\%$)	Script (Latin or not)	Lang. family (Indo-European or not)
XGLM-7.5b	Pretrain data ($\leq 0.003\%$)	Pretrain data ($\leq 0.15\%$)	Lang. family (Indo-European or not)
Full-shot			
Model	ALL	SEEN	UNSEEN
mBERT	Pretrain data ($\leq 0.032\%$)	Pretrain data ($\leq 0.073\%$)	Lang. family (Indo-European or not)
XLNet	Pretrain data ($\leq 0.005\%$)	Pretrain data ($\leq 0.031\%$)	Lang. family (Indo-European or not)

Table 2: Top features identified by decision tree analysis for each model and scenario. For SEEN languages, pretraining data size and resource level dominate (except for XGLM-2.9b, where script type is most influential). For UNSEEN languages, linguistic characteristics (script type and language family) take precedence. All features exhibit statistically significant differences in performance ($p < 0.001$).

ALL Languages:

- For the ALL languages scenario, decision trees clearly reveal that pretraining data is the most influential factor in 29 out of 31 cases. Because ALL includes languages SEEN and UNSEEN,

notably, our deeper look at the decision tree analyses indicates that this factor in most cases boils down to *whether the language was part of the training set or not, rather than the amount of language-specific data*, as indicated by the values

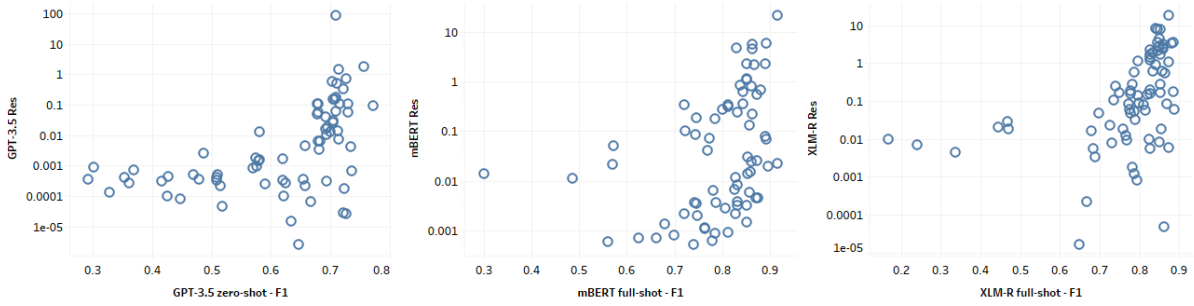


Figure 3: F1 Score vs. model-specific pretraining data (percentage) for GPT-3.5, mBERT and XLM-R models.

of the pretraining data percentages which range from 0% to at most 0.175%. GPT-3.5 model draws the distinction along general resource levels whether a language is low resource (0, 1, or 2) or level 3 and higher.

SEEN Languages:

- For SEEN languages, model-specific pretraining data continues to remain the most influential factor in 22 out of 31 model and scenario combinations. However, this time because there are no unseen languages in the mix, the model performance seems to be impacted by the amount of pretraining data, as indicated by the slightly higher percentage values as compared to the ALL languages scenario.
- Interestingly, general resource availability based on linguistic diversity taxonomy (Joshi et al., 2020) appears to be the most important factor for Bloom models in the zero-shot setup, as well as for xglm-564m (zero-shot) and xglm-1.7b (two-shot). For Bloom models, the distinction is along resource levels 0/1/2 or higher, whereas for xglm models, it is along 0/1 and higher. Additionally, xglm-2.9b in both zero-shot and two-shot scenarios shows a stronger influence of script type (Latin or not). These cases indicate that factors beyond pretraining data size can also play a significant role for specific models and settings.
- Furthermore, Figure 3 plots the performance of mBERT, XLM-R, and GPT-3.5 models in relation to model-specific pretraining data amounts. The figure demonstrates a clear trend: as the model-specific language data increases, so does the model’s performance. This observation aligns with the finding that pretraining data size is a crucial factor for SEEN languages.

UNSEEN Languages:

- In contrast to SEEN languages, UNSEEN languages show quite a different pattern. Naturally, because UNSEEN languages do not have pretraining data as one of their relevant factors, it is absent from this column. However, out of 31 models, 23 are most impacted by script type, and 8 are most influenced by language family. This shift in importance towards linguistic features suggests that when models encounter unfamiliar languages, they rely more heavily on similarities in writing systems to generalize from their existing knowledge.
- Within the scripts and language families, there are nuanced differences. For instance, while generally the models make the distinction along the lines of whether the script is Latin or not, occasionally Devanagari script also seems important, particularly for XGLM models. Similarly, while Indo-European is the most common influential language family, we also observe an instance each of Austronesian and Niger-Congo. Additionally, models of different sizes from the same family may prefer not just a different script or a different language family when moving from zero-shot to two-shot setting, they may prefer an entirely different factor (e.g., Bloom-560m in zero-shot vs. two-shot settings), further complicating the matters.

5 Discussion

Our comprehensive analysis of 6 multilingual models on the SIB-200 dataset reveals valuable insights into the factors influencing their performance across a diverse range of languages.

Our key findings can be summarized as follows:

- Pretraining data size consistently emerges as a crucial factor, but the distinction is less along

the quantity of data but rather whether the languages have been encountered during training or not.

- For UNSEEN languages, script type and language family are influential, suggesting that MLLMs rely on cross-lingual transfer learning to generalize to unfamiliar languages.
- General resource availability plays a less prominent role overall but appears to be important for one specific model under one setting (Bloom in zero-shot for seen languages).
- Interestingly, the performance of Bloomz, an instruction-tuned model, is more influenced by the distribution of languages in its pretraining corpus than the fine-tuned dataset used for instruction tuning. This suggests that the initial pretraining stage plays a crucial role in shaping the model’s capabilities, even after further fine-tuning for specific tasks.
- Finally, our analysis also indicates that while model size and architecture may influence overall performance, they do not significantly alter the most important features identified by the decision trees. The distribution of languages in the pretraining data and the linguistic characteristics of the target languages consistently emerge as the dominant factors regardless of the specific model architecture or scale.

Several future directions remain to be explored. We observed that script type can be more influential for specific models and settings. Further investigation is needed to understand the reasons behind these preferences and how they can be leveraged to achieve more consistent performance across languages. It is also not clear why models lean towards different factors under different settings (for instance, resource level is important in Bloom-560m zero-shot setting but pretraining data is important in its two-shot ICL setting).

6 Conclusion

This study analyzed 6 multilingual language models on the SIB-200 dataset, revealing key insights into their performance across around 200 languages. We found that the size of the pretraining data significantly affects performance. For unseen languages, script type and language family become

more crucial, highlighting the importance of cross-lingual transfer learning. While general resource availability plays a less prominent role overall, it can be significant for specific models and settings. Interestingly, model size and architecture do not significantly change the most important features identified in our analysis. Our work contributes to a deeper understanding of MLLMs and hopes to guide the development of more effective and equitable multilingual NLP systems.

Limitations

This study provides insights into multilingual language model performance, but it is important to acknowledge certain limitations. The SIB-200 dataset, while extensive, may contain biases in language representation and genre distribution, potentially affecting the generalizability of our findings. Additionally, our analysis focuses on the text classification task, and the findings may not directly generalize to other NLP tasks. While we analyzed a diverse set of models, our findings may not be fully representative of the entire MLLM landscape. Finally, our analysis is based on the current state of MLLMs, and the relative importance of different factors may change as these models continue to evolve. Future research should address these limitations by expanding to more diverse datasets, investigating different NLP tasks, evaluating a broader range of models, and conducting longitudinal studies.

Ethics Statement

The experimental setup and code implementation ensured adherence to ethical guidelines, data usage agreements, and compliance with the terms of service of the respective language models and data sources. The research team also recognized the importance of inclusivity and fairness by considering a diverse set of languages and language families in the evaluation, thus avoiding biases and promoting balanced representation.

Acknowledgements

We are grateful to the anonymous reviewers whose feedback and thought-provoking questions enhanced this paper. The engaging discussions and collaborative spirit within the PortNLP research group were instrumental in shaping this research. We acknowledge the National Science Foundation for their financial support through grants (CRII:RI

2246174 and SAI-P 2228783), which made this work possible.

References

- David Ifeoluwa Adelani, Hannah Liu, Xiaoyu Shen, Nikita Vassilyev, Jesujoba O Alabi, Yanke Mao, Haonan Gao, and Annie En-Shiun Lee. 2023. Sib-200: A simple, inclusive, and big evaluation dataset for topic classification in 200+ languages and dialects. *arXiv preprint arXiv:2309.07445*.
- Ameeta Agrawal, Lisa Singh, Elizabeth Jacobs, Yaguang Liu, Gwyneth Dunlevy, Rhitabrat Pokharel, and Varun Uppala. 2023. [All translation tools are not equal: Investigating the quality of language translation for forced migration](#). In *2023 IEEE 10th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 1–10.
- Kabir Ahuja, Rishav Hada, Millicent Ochieng, Prachi Jain, Harshita Diddee, Samuel Maina, Tanuja Ganu, Sameer Segal, Maxamed Axmed, Kalika Bali, et al. 2023. Mega: Multilingual evaluation of generative ai. *arXiv preprint arXiv:2303.12528*.
- Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepy, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, YaGuang Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni, Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussalem, Zachary Nado, John Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Aurko Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R. So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhong Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, and Yonghui Wu. 2023. [Palm 2 technical report](#).
- Francesco Barbieri, Luis Espinosa Anke, and Jose Camacho-Collados. 2022. [XLM-T: Multilingual language models in Twitter for sentiment analysis and beyond](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 258–266, Marseille, France. European Language Resources Association.
- Damian Blasi, Antonios Anastasopoulos, and Graham Neubig. 2022. [Systematic inequalities in language technology performance across the world’s languages](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5486–5505, Dublin, Ireland. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Tyler A Chang and Benjamin K Bergen. 2023. Language model behavior: A comprehensive survey. *arXiv preprint arXiv:2303.11504*.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pilla, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhong Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [Palm: Scaling language modeling with pathways](#).
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ethnologue. 2022. [What are the largest language families?](#)
- Fahim Faisal, Orevaoghene Ahia, Aarohi Srivastava, Kabir Ahuja, David Chiang, Yulia Tsvetkov, and Antonios Anastasopoulos. 2024. [Dialectbench: A nlp benchmark for dialects, varieties, and closely-related languages](#).
- Ayyoob ImaniGooghari, Peiqin Lin, Amir Hossein Kargar, Silvia Severini, Masoud Jalili Sabet, Nora Kassner, Chunlan Ma, Helmut Schmid, André Martins, François Yvon, and Hinrich Schütze. 2023. [Glot500: Scaling multilingual corpora and language models to 500 languages](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1082–1117, Toronto, Canada. Association for Computational Linguistics.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The state and fate of linguistic diversity and inclusion in the NLP world](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Viet Dac Lai, Nghia Trung Ngo, Amir Pouran Ben Veyseh, Hieu Man, Franck Dernoncourt, Trung Bui, and Thien Huu Nguyen. 2023. [Chatgpt beyond english: Towards a comprehensive evaluation of large language models in multilingual learning](#). *arXiv preprint arXiv:2304.05613*.
- 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 Diab, Veselin Stoyanov, and Xian Li. 2022. [Few-shot learning with multilingual language models](#).
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual Denoising Pre-training for Neural Machine Translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Henry B Mann and Donald R Whitney. 1947. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, pages 50–60.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 2023. [Crosslingual generalization through multitask finetuning](#).
- OpenAI. 2023. [Gpt-4 technical report](#).
- Fred Philippy, Siwen Guo, and Shohreh Haddadan. 2023. [Towards a common understanding of contributing factors for cross-lingual transfer in multilingual language models: A review](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5877–5891, Toronto, Canada. Association for Computational Linguistics.
- Surangika Ranathunga and Nisansa de Silva. 2022. [Some languages are more equal than others: Probing deeper into the linguistic disparity in the nlp world](#). *arXiv preprint arXiv:2210.08523*.
- Teven Le Scao, Angela Fan, Christopher Akiki, Ellice Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2022. [Bloom: A 176b-parameter open-access multilingual language model](#). *arXiv preprint arXiv:2211.05100*.
- Oleh Shliakhko, Alena Fenogenova, Maria Tikhonova, Vladislav Mikhailov, Anastasia Kozlova, and Tatiana Shavrina. 2022. [mgpt: Few-shot learners go multilingual](#). *arXiv preprint arXiv:2204.07580*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#).
- Don Vaughan. 2020. [The world’s 5 most commonly used writing systems](#).
- Wenxuan Wang, Zhaopeng Tu, Chang Chen, Youliang Yuan, Jen-tse Huang, Wenxiang Jiao, and Michael R Lyu. 2023. [All languages matter: On the multilingual safety of large language models](#). *arXiv preprint arXiv:2310.00905*.
- Shijie Wu and Mark Dredze. 2020. [Are all languages created equal in multilingual bert?](#) *arXiv preprint arXiv:2005.09093*.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

Wen Yang, Chong Li, Jiajun Zhang, and Chengqing Zong. 2023. [Bigtrans: Augmenting large language models with multilingual translation capability over 100 languages](#). *arXiv preprint arXiv:2305.18098*.

Wenhao Zhu, Hongyi Liu, Qingxiu Dong, Jingjing Xu, Lingpeng Kong, Jiajun Chen, Lei Li, and Shujian Huang. 2023. Multilingual machine translation with large language models: Empirical results and analysis. *arXiv preprint arXiv:2304.04675*.

A Appendix: Prompts

This appendix provides the specific prompts used for evaluating Bloom, Bloomz, XGLM, and GPT-3.5 in the zero-shot and two-shot in-context learning (ICL) settings on the SIB-200 text classification task.

Zero-shot Prompt (Bloom, Bloomz, XGLM):

```
SENTENCE: "{input sentence}"
Is this SENTENCE science, travel, politics,
sports, health, entertainment, geography?
OPTIONS:
-science
-travel
-politics
-sports
-health
-entertainment
-geography
ANSWER:
```

Two-shot ICL Prompt (Bloom, Bloomz, XGLM):

What category does SENTENCE belong to?

```
SENTENCE: "{sentence1}"
LABEL: {label1}
SENTENCE: "{sentence2}"
LABEL: {label2}
...
SENTENCE: "{sentence14}"
LABEL: {label14}
SENTENCE: "{input sentence}"
OPTIONS:
-science
-travel
-politics
-sports
-health
-entertainment
-geography

LABEL:
```

Zero-shot Prompt (GPT-3.5):

You will be provided with a text, and your task is to classify its category as science, travel, politics, sports, health, entertainment, geography.
{input sentence}

Category:

B Appendix: Supplemental plots

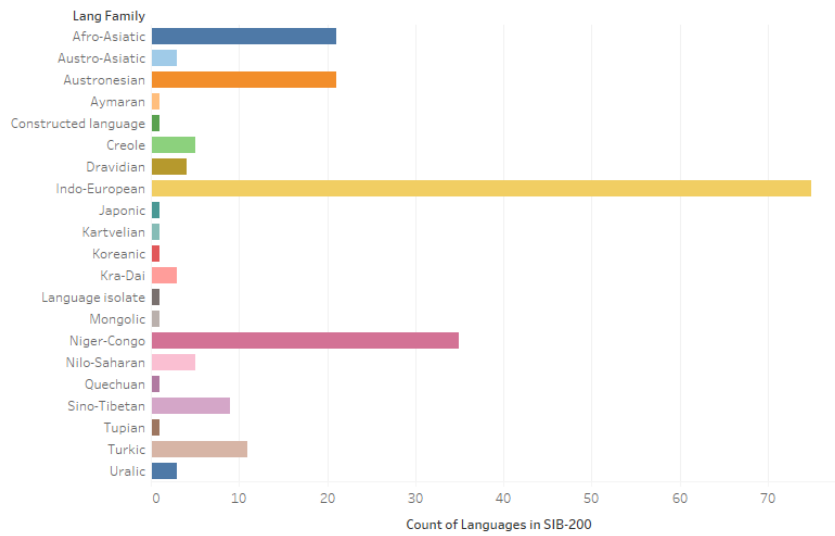


Figure 4: Distribution of language family in SIB-200.

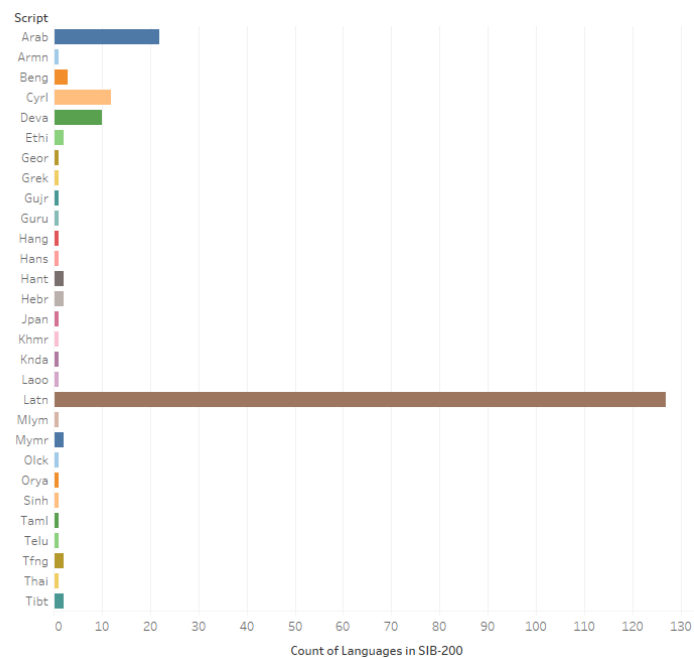


Figure 5: Distribution of scripts in SIB-200.