

MC²: Towards Transparent and Culturally-Aware NLP for Minority Languages in China

Chen Zhang*, Mingxu Tao*, Quzhe Huang*,
Jiuheng Lin*, Zhibin Chen, Yansong Feng†

Peking University

{zhangch, thomastao, huangquzhe, czb-peking, fengyansong}@pku.edu.cn
linjiuheng@stu.pku.edu.cn

Abstract

Current large language models demonstrate deficiencies in understanding low-resource languages, particularly the minority languages in China. This limitation stems from the scarcity of available pre-training data. To address this accessibility challenge, we present MC², a **Multilingual Corpus of Minority Languages in China**, which is the largest open-source corpus of its kind so far. MC² includes four underrepresented languages: Tibetan, Uyghur, Kazakh, and Mongolian. Notably, we focus on the less common writing systems of Kazakh and Mongolian, i.e., Kazakh Arabic script and traditional Mongolian script, respectively, which have been long neglected in previous corpus construction efforts. Recognizing the prevalence of language contamination within existing corpora, we adopt a quality-centric solution for collecting MC², prioritizing accuracy while enhancing diversity. Furthermore, we underscore the importance of attending to the multiplicity of writing systems, which is closely related to the cultural awareness of the resulting models. The MC² corpus and related models are made public to the community¹.

1 Introduction

Recently, the rapid development of large language models (LLMs) has been fueled by the availability of high-quality pre-training data. However, only a handful of high-resource languages, such as English and Chinese, have benefitted from the advantages of this progress in LLMs. Despite having a substantial user base, many languages remain excluded from the benefits of these advancements due to a lack of suitable corpora. In this paper, we focus on such underrepresented languages, particularly addressing the minority languages in China, including Tibetan, Uyghur, Kazakh, and Mongolian.

*Equal contribution.

†Corresponding author.

¹https://github.com/luciusssss/mc2_corpus




Figure 1: The population of native speakers vs. the number of tokens collected by CulturaX, a representative multilingual corpus. Green marks represent the four Chinese minority languages studied in this work, with orange marks for several mid-resource languages in Asia and violet marks for six languages with the richest resources. We obtain the population data from Wikipedia and check the credibility of the references.

As illustrated in Figure 1, the four languages, although spoken by tens of millions, face a critical deficiency in linguistic resources, posing obstacles to both academic research and the development of practical AI applications. Notably, there is no open-source corpus available for Kazakh in Arabic script and Mongolian in traditional Mongolian script – two language variants in China adopting underrepresented writing systems. In the case of Tibetan and Uyghur, although several multilingual datasets such as OSCAR (Abadji et al., 2022) and CulturaX (Nguyen et al., 2023) cover these languages, their quality falls short. Our preliminary study reveals a critical issue of language misidentification. Up to 34% of data in the Uyghur split of CulturaX is actually Kazakh or Arabic texts. This raises concerns for future research in low-resource NLP, as conducting experiments on these significantly contaminated *monolingual* corpora may yield misleading conclusions.

Name	ISO 639-1	ISO 639-3	Language Family	Writing System
Tibetan	bo	bod	Sino-Tibetan	Tibetan script
Uyghur	ug	uig	Turkic	Uyghur Arabic script
Kazakh	kk	kaz	Turkic	Kazakh Arabic script
Mongolian	mn	mvf	Mongolic	Traditional Mongolian script

Table 1: ISO codes, languages families, and writing systems of the languages in MC².

	CC100	mC4	ROOTS	OSCAR	CulturaX	MC ² (crawl)	MC ² (full)
Tibetan	✗	✗	✗	532M	718M	1.7G	2.2G
Uyghur	378M	✗	✗	146M (220M) [†]	338M (403M) [†]	520M	736M
Kazakh (Arabic)	✗	✗	✗	✗	✗	937M	937M
Mongolian (Traditional)	✗	✗	✗	✗	✗	970M	970M

Table 2: Dataset sizes of different multilingual datasets in the four minority languages. MC² (crawl) denotes the subset of our newly-collected web crawls. MC² (full) is the complete set of MC², which additionally contains texts collected from existing resources. [†]We observe many cases of language misidentification in the Uyghur split of OSCAR and CulturaX. We report the data sizes after manual re-identification. The crossed-out numbers in the brackets indicate the original data sizes.

To facilitate more transparent and reproducible NLP research on the minority languages in China, we present MC², a **Multilingual Corpus of Minority Languages in China**, which is the largest open-source corpus for these languages so far. During data collection, we adhere to a quality-centric principle. We carefully design strategies for the selection of web pages to crawl, ensuring the language purity of the crawled texts. Afterwards, we conduct a thorough data cleaning process, eliminating extraneous elements such as web page headers and retaining only meaningful content. Our newly collected MC² dataset surpasses existing datasets by containing more than twice the amount of Tibetan and Uyghur data, while maintaining superior data quality.

MC² particularly emphasize underrepresented writing systems, recognizing their significance both technically and culturally. In languages with multiple writing systems, such as Kazakh and Mongolian, previous research has predominantly focused on the more common scripts (Abadji et al., 2022; Nguyen et al., 2023). We illustrate that transliterating texts from high-resource scripts into lower-resource ones may not be effective for training, especially when the script conversion is flawed. Moreover, the texts collected in different scripts inherently embody cultural nuances of various language communities. Through probing, we demonstrate that writing systems play a crucial role in developing culturally-aware NLP systems.

In addition, we continually train language models with MC². They achieve comparable perfor-

mance to their counterparts trained with closed-source data. This further validates the value of our data in facilitating transparent NLP research.

To summarize, we make the following contributions: (1) We present MC², the largest open-source corpus to date for four underrepresented languages in China. (2) We highlight quality issues, especially language misidentification, prevalent in previous corpora, which may threaten trustworthy research on low-resource languages. (3) We reveal the cultural complexities arising from the multiplicity of writing systems within languages.

2 Related Works

Multilingual Corpus To facilitate the development of multilingual LLMs, a series of large-scale multilingual corpora has been released, including CC100 (Conneau et al., 2020), mC4 (Raffel et al., 2020), ROOTS (Laurençon et al., 2022), OSCAR (Abadji et al., 2022), CulturaX (Nguyen et al., 2023), and Madlad-400 (Kudugunta et al., 2024). These corpora cover both high-resource languages and a handful of low-resource ones. Recently, there are also efforts in constructing open-source corpora for low-resource languages in certain regions (Cahyawijaya et al., 2021; Teodorescu et al., 2022; Doddapaneni et al., 2023). However, the minority languages in China are still underrepresented in existing datasets, which calls for more attention to constructing open-source corpus for these languages.

NLP for Minority Languages in China In terms of the minority languages in China, previous works try to improve the accessibility of these languages, by collecting a handful of annotated datasets for specific NLP tasks, each covering one or some languages investigated in our study. They mainly focus on three types of tasks: text classification (Qun et al., 2017; Yang et al., 2022; Shi et al., 2022b), question answering (Sun et al., 2021), and machine translation (Costa-jussà et al., 2022; Zhang et al., 2024). The models trained specifically for the minority languages in China include CINO (Yang et al., 2022), MiLMo (Shi et al., 2022b), and CMPT (Li et al., 2022). However, none of these works releases their pre-training corpus. In contrast, we release a large-scale high-quality corpus for four minority languages in China, which, we hope, will help relieve the data scarcity problem and improve transparency in research on low-resource languages.

Culturally-Aware NLP Cultural considerations are increasingly recognized as crucial in NLP research and applications (Hershcovich et al., 2022). An important aspect is the connection between cultural factors and linguistic forms within languages. Previous works on intra-language variations have primarily focused on dialects (Zampieri et al., 2020; Ziemis et al., 2022; Liu et al., 2023) and sociolects (McCormack et al., 2011; Zhang et al., 2021). However, the dimension of writing systems within languages has received little attention. We demonstrate that texts in underrepresented writing systems may encode unique cultural nuances that differ significantly from those found in communities using more common scripts.

3 Data Collection and Analysis

Here we describe the procedure of creating MC², the largest open-source corpus so far for four minority languages in China.

Before collecting MC², we first conduct a preliminary audit in the low-resource language splits of previous multilingual corpora and find severe quality issues such as language contamination. Therefore, we propose a quality-centric solution for data collection of low-resource languages, which aims to ensure accuracy while improving the comprehensiveness and coverage of the data.

The collection procedure of MC² mainly consists of three steps: (1) gathering web crawls, (2) incorporating existing datasets, and (3) deduplicat-

ing and filtering. Throughout the collection process, we adhere to the protocol of quality-centric collection. We hope to establish a reliable groundwork for subsequent language model training or linguistic research. Notably, our framework for data collection is generally language-agnostic and can be easily applied to collecting web corpora for other low-resource languages.

3.1 Quality Issues in Previous Corpora

When auditing previous multilingual web corpora for low-resource languages (Abadji et al., 2022; Nguyen et al., 2023), we find two critical quality issues: language misidentification and insufficient data cleaning. These defects pose a significant threat to effective model training and might undermine the credibility of research findings.

Language Misidentification An important step in collecting web corpora for low-resource languages is to identify the language of a web page. However, current language identification tools such as fastText (Joulin et al., 2016) are prone to error (Kreutzer et al., 2022). This issue is more severe in languages with similar writing systems.

Specifically, we audit the Uyghur splits in CulturaX and OSCAR. Since the documents in these two datasets are paired with their respective URLs, we can efficiently label their languages according to the website of their URLs. We manually check the 654 websites comprising the Uyghur corpus in CulturaX. Our examination reveals that 16% of the data is actually in Kazakh or Arabic, languages utilizing scripts akin to that of Uyghur. Similarly, 33% of the Uyghur corpus in OSCAR are Kazakh or Arabic. See the example of misidentification in Appendix A.1.

This issue presents a significant concern for future research in low-resource NLP, as experiments conducted on heavily contaminated corpora can produce misleading results. For instance, this contamination can adversely impact research on language transfer. Researchers may aim to evaluate whether a model trained on a *monolingual* corpus in Language A can transfer its capabilities to Language B. However, if the corpus already contains a considerable amount of Language B, the target language, the conclusions from such experiments are unlikely to be reliable.

Insufficient Data Cleaning The web crawls in previous corpora often contain unwanted texts such

as sidebars, headers, and footers, as shown the example in Appendix A.1. We sample 100 pages from the Tibetan corpus of CulturaX and our manual annotation shows that 42% contain headers or footers. These undesired texts affect the coherence of the document texts, which might hinder models from learning the linguistic patterns of low-resource languages.

3.2 Step 1: Web Crawling

Our corpus is mainly made up of web crawls. We combine both human labor and AI assistance to prevent the flaws in the previous corpora.

Language Identification Different from previous efforts for curating *massively-multilingual* web crawls (Nguyen et al., 2023), the number of websites in each language is limited for the four minority languages in our study. We manually maintain a list of high-quality websites for each language of our study, to avoid language contamination resulting from mislabeling by identification tools. These websites are collected from a variety of sources, including URLs in existing corpora, search engines, web portals, website recommendations on social media, and hyperlinks on these websites. For each website in the list, we start from the homepage and crawl all the pages through a breadth-first search.

Text Extraction The web crawls in previous corpora often contain unwanted texts such as sidebars, headers, and footers. As the number of websites in our corpus is manageable and the web pages in the same website usually share the same structure, we automatically design text extraction rules for each website. These rules can precisely extract the title and main content of a web page, discarding other distracting texts. Specifically, for each website, we ask Github Copilot² to analyze the HTML structure of a sampled web page and write a Python code to extract its title and main content. In this way, we can extract the wanted texts in the raw web crawls with high accuracy and efficiency.

In total, we collected 2.0G Tibetan, 1.1G Uyghur, 1.3G Kazakh, and 1.7G Mongolian texts in our new web crawls.

3.3 Step 2: Incorporation of Existing Datasets

Thanks to the existing efforts in the community, we incorporate open-source resources into our corpus, including **CulturaX** (Nguyen et al., 2023),

²<https://github.com/features/copilot>

Wikipedia³, and **NLGIW 2023 Shared Task**⁴. We remove the data with language contamination when merging these datasets.

CulturaX is a cleaned collection of mC4 and OSCAR. Its Tibetan and Uyghur splits can be used to supplement our corpus. However, there are errors in language identification in CulturaX, especially the confusion between Uyghur texts and Kazakh texts in Arabic scripts. We thus conduct language re-identification on the Uyghur split of CulturaX, using the website-centric strategy explained in Section 3.1. From CulturaX, we obtained 718M Tibetan, 338M Uyghur, and 65M Kazakh texts.

Wikipedia is a high-quality resource for low-resource languages. We collected 127M Tibetan and 41M Uyghur texts from Wikipedia. NLGIW 2023 Shared Task provides an open-source corpus of Tibetan news, where we add 342M texts.

3.4 Step 3: Deduplication and Filtering

Previous works claim that properly filtered and deduplicated web data is crucial to effective pre-training (Lee et al., 2022; Penedo et al., 2023; Ranathunga et al., 2024). Therefore, we take a series of measures for deduplication and filtering to ensure the high quality of our corpus.

Deduplication We find two sources of duplication in our corpus. One is that the websites in minority languages often repost from each other. The other is the duplication between the web crawls in Step 1 and the existing corpus merged in Step 2. We thus use a combination of deduplication methods, including URL-based, exact, and fuzzy deduplication. See the implementation details in Appendix A.2

Filtering We design the filtering rules from several aspects, including repetition, document lengths, and unexpected characters. See details in Appendix A.2.

We additionally make efforts to remove privacy information from the corpus. We use heuristics to identify emails, telephone numbers, and Resident Identity Card numbers. We replace them with special tokens, i.e., [email], [phone], and [idcard].

After deduplication and filtering, we obtain 2.2G Tibetan, 736M Uyghur, 937M Kazakh, and 970M Mongolian texts, which compose the current version of MC².

³<https://huggingface.co/datasets/graelo/wikipedia>

⁴http://nlg.cipsc.org.cn/evaluation_12.html




Figure 2: Domain distributions of MC². The percentages are calculated from data sizes.

3.5 Statistical Analysis of MC²

We provide an in-depth analysis from the perspective of domain diversity and document length.

Domain Diversity For the pretraining corpus, both quality and diversity are crucial. When curating the website list for crawling, we comprehensively collect high-quality websites from diverse domains to enhance the dataset’s diversity while ensuring quality. Specifically, our data sources fall into five categories: News, Culture, Government, Wikipedia, and Others. Figure 2 shows the proportion of different categories of websites in MC². Our corpus covers formal official documents, and miscellaneous news articles about politics, important events, daily life, and entertainment, as well as cultural content related to literature, art, and religion. It is noteworthy that news articles constitute the majority of our datasets. News articles are usually grammatically correct, which could potentially facilitate learning language modeling under a low-resource setting.

Document Length Constrained by the absence of high-quality long-form texts, prior efforts like CINO (Yang et al., 2022) have to train their models with a collection of individual sentences, which might restrict their abilities on longer texts. In contrast, our work introduces a document-level corpus. We report the distribution of document length for each language in Figure 3. The average document length of MC² far exceeds the maximum input length, i.e., 512, of the previous BERT-size model for minority languages in China.

We observe that CultureX also includes long documents. However, as mentioned in Section 3.2, its data cleaning might not be adequate, introducing




Figure 3: Length distributions of the documents in MC² (using log scales for x-axes), calculated based on the tokens by the tokenizer of CINO (Yang et al., 2022).

Script	Region	Sample Text	Size
Arabic (kk)	China	سالهه تىسىز به	0
Cyrillic (kk)	Kazakhstan	сәлеметсіз бе	18G
Traditional (mn)	China	ᠰᠠᠶᠢᠨ ᠤᠮᠤ	0
Cyrillic (mn)	Mongolia	сайн уу	12G

Table 3: Comparison between the different writing systems of Kazakh (kk) and Mongolian (mn). The sample texts mean *hello*. We report the data sizes in CulturaX.

undesired texts like sidebars, headers, and footers. These issues could disrupt the semantic coherence of the entire document, making it challenging for models to learn correct dependencies within long texts. We expect that our high-quality, long-form corpus will assist future efforts in training models with better capacities for understanding and generating longer texts in minority languages.

4 Discussion on Underrepresented Writing Systems

Many languages adopt distinct writing systems across various regions. For instance, in China, minority languages such as Kazakh and Mongolian employ scripts that differ from the Cyrillic scripts used in Kazakhstan and Mongolia, as illustrated in Table 3. Unfortunately, existing datasets predominantly concentrate on the more prevalent writing systems, neglecting the less common ones. In re-



Figure 4: Comparison between training XLM-RoBERTa-large with authentic and transliterated traditional Mongolian data. (a) The loss on the training data. (b) The perplexity on the hold-out evaluation data from the authentic traditional Mongolian corpus.

sponse to this issue, MC² is the first effort to collect native corpora for the two underrepresented writing systems, i.e., the Kazakh Arabic script and the traditional Mongolian script.

To delve into the multiplicity of writing systems, we conduct empirical studies from both the **technical** and the **cultural** perspectives. We find that imperfect transliteration between writing systems harms model training. Furthermore, corpora in different writing systems of the same language encapsulate unique cultural context specific to their respective language communities. These findings underscore the importance of gathering resources for minority writing systems.

4.1 Flawed Training with Imperfect Transliteration

A common approach to addressing different writing systems is transliteration (Nakov and Ng, 2009; Chakravarthi et al., 2019; Muller et al., 2021; Sun et al., 2022). To obtain a model for low-resource scripts, it is intuitive to transliterate the corpus in the high-resource scripts into low-resource ones for training. However, there are no one-to-one conversion rules between scripts for languages such

as Mongolian (Bao et al., 2014). The transliteration between traditional and Cyrillic Mongolian is context-dependent and current open-source tools are far from perfect. We investigate how the noisy transliteration results affect the training of models.

Experimental Setup We prepare two versions of training data in traditional Mongolian. One is 900M *transliterated* traditional Mongolian data, which is converted from the Cyrillic Mongolian data in CulturaX, using the only open-source transliteration tool available so far⁵. The transliterated results often contain errors, as shown by the examples in Appendix C. The other version of data is an equivalent volume of *authentic* traditional Mongolian data in MC². We then train XLM-RoBERTa-large (Conneau et al., 2019) with these two different versions of data respectively and compare the convergence speed and performance of the models.

Results The training processes are illustrated in Figure 4. The training loss of transliterated data is significantly higher than that of authentic data. This discrepancy arises due to the substantial noise in transliterated data, making it challenging for the model to capture underlying linguistic patterns. We further measure the perplexity of both models on a hold-out validation set of traditional Mongolian text. The perplexity of the model trained on authentic data steadily decreases with the training steps, maintaining a similar magnitude to the training loss. Conversely, the model trained on transliterated data exhibits a sharp increase in perplexity when evaluated on authentic data, indicating that it fails to learn the language modeling of true traditional Mongolian.

Lessons Learned Using noisy data transliterated from high-resource writing systems will greatly hinder the learning of low-resource writing systems. This highlights the necessity of collecting an authentic corpus for underrepresented scripts in the absence of mature transliteration tools.

4.2 Cultural Differences Behind Writing Systems

For some languages such as Kazakh, we can achieve perfect transliteration between different writing systems using pre-defined rules. Nevertheless, there exist disparities in the cultural backgrounds between the language variants using dif-

⁵<https://github.com/tugstugi/mongolian-nlp>

ferent scripts. With the technique of probing (Jiang et al., 2020; Zhang et al., 2021), we investigate whether the training data collected from different writing systems will lead to distinct cultural knowledge in the resulting models.

Probing Questions We take the Kazakh language as our research target. The Kazakh community in China uses the Arabic script while the Cyrillic script is adopted in Kazakhstan. We design probing questions that reflect the cultural differences between the two Kazakh communities, such as those in Table 4. Each question offers two options: Option A aligns with the cultural context of the Kazakh community in China, while Option B corresponds to the cultural context of Kazakhstan. These questions are manually formulated by comparing Wikipedia articles that describe the two communities. We compile 25 pairs of questions covering topics such as geography, politics, and daily life. Originally written in English, these questions are translated into Cyrillic Kazakh using Google Translate, and then transliterated into Arabic Kazakh.

Experimental Setup We train two distinct Kazakh language models based on XLM-RoBERTa-large, each tailored to one of the writing systems. One is trained with 900M authentic Cyrillic Kazakh texts from CulturaX. And the other is trained with an equivalent volume of Arabic Kazakh texts from our MC² corpus. We subject the two models trained on different scripts to the cultural probing questions. We query the Arabic Kazakh model with questions written in the Arabic script. Similarly, for the Cyrillic Kazakh model, we use questions written in the Cyrillic script.

Results The Arabic Kazakh model selects Option A for 84% of the probing questions, indicating a strong alignment with the cultural characteristics of Kazakh communities in China. Conversely, the Cyrillic Kazakh model chooses Option B for 56% of the questions, reflecting the cultural characteristics of Kazakh communities in Kazakhstan.

The cultural distinction is probably inherent in the corpus collected from websites using different writing systems. Subsequently, it is captured by the models trained on data in different scripts.

In Table 4, we present three examples of probing questions. For Query 1, the two models yield divergent responses concerning the holiday celebrated on May 1st. This discrepancy arises from the fact that the official holidays on this date differ

Option	Arabic	Cyrillic
<i>Query 1: On May 1st we celebrate ____.</i>		
A. International Workers' Day	2.48	5.46
B. People's Unity Day	3.23	3.30
<i>Query 2: The currency used here is ____.</i>		
A. Renminbi	1.97	7.29
B. Kazakhstani Tenge	3.33	6.55
<i>Query 3: There is a big lake near my home called ____.</i>		
A. Sayram Lake	3.33	3.15
B. Lake Balkhash	4.35	2.66

Table 4: Knowledge probing on the two models trained with Kazakh data in different writing systems. The numbers are the perplexities of the corresponding options.

between China and Kazakhstan. Similarly, Query 2 and Query 3 reflect the cultural differences in terms of economy and geography. See their explanations in Appendix B.

Lessons Learned The corpora in different writing systems contain unique cultural context specific to their respective communities. Hence, it is sub-optimal to directly transliterate high-resource scripts into low-resource ones for training, which may undermine the cultural nuances underlying the low-resource writing system. The inclusiveness of writing systems is crucial to the construction of culturally-aware models. We encourage future research on properly leveraging the data in the high-resource writing systems while preserving the cultural uniqueness in the underrepresented scripts.

5 Continual Pretraining with MC²

To demonstrate the practical value of our corpus, we train models with MC² and compare their performance with competitive counterparts. Instead of training from scratch, we conduct continual pretraining on existing models with MC², which is an effective technique to adapt models to new languages (Ebrahimi and Kann, 2021; Muller et al., 2021; Yong et al., 2023). We continually pretrain XLM-RoBERTa-large (Conneau et al., 2019) with MC², obtaining **MC²XLMR-large**. Thanks to the high quality of MC², it performs comparably to CINO, a counterpart trained with closed-source data. We also attempt to adapt larger models such as Llama (Touvron et al., 2023) to the four languages, obtaining **MC²Llama-13B**. The newly trained model shows the superior ability of in-context learning in these languages.

Model	Open-source Corpus	Classification (Accuracy / Weighted F1)					QA (EM / F1)
		bo	ug	kk	mn	All	bo
mBERT-base	No	20.9 / 19.8	67.7 / 77.4	13.7 / 10.6	56.1 / 59.6	27.8 / 26.0	3.6 / 4.6
XLM-RoBERTa-large	No	32.3 / 28.8	86.3 / 88.7	35.8 / 34.8	43.5 / 45.8	39.0 / 38.0	7.0 / 30.1
CINO-large-v2	No	55.9 / 58.6	88.0 / 89.5	42.0 / 41.6	<u>62.4 / 65.2</u>	<u>50.4 / 50.1</u>	<u>14.0 / 57.1</u>
MC ² XLMR-large (Ours)	Yes	<u>52.6 / 49.9</u>	87.7 / 89.7	<u>40.3 / 37.4</u>	72.6 / 74.4	52.0 / 48.9	14.8 / 53.6

Table 5: Performance of different models under the zero-shot transfer setting. The best scores are made **bold**, with the second underlined. For classification, we use accuracy and weighted F1 as metrics. For QA, we use EM and F1.

5.1 Model Training

Two models with different architectures and sizes are used for continual pretraining. One is XLM-RoBERTa-large (Conneau et al., 2019), a 560M encoder-only multilingual model. The other is Llama2-ZH-13B (Huang et al., 2023), a bilingual model supporting both English and Chinese.

The vocabularies of XLM-RoBERTa-large and Llama2-ZH-13B hardly contain tokens for the four languages in our study. Thus, we add 3K new tokens for each language, obtained by BPE (Sennrich et al., 2016).

To avoid the catastrophic forgetting of learned languages, we add 0.25B tokens of Chinese data from Wanjuan (He et al., 2023) and 0.25B tokens of English data from C4 (Raffel et al., 2020) for training. See more training details in Appendix D.1.

5.2 Evaluation Setup

Datasets For the four minority languages, we can only find limited datasets for evaluation. **WCM-v2** (Yang et al., 2022) is a 10-category text classification dataset. **TibetanQA** (Sun et al., 2021) is a Tibetan machine reading comprehension (MRC) dataset. **Flores-200** (Costa-jussà et al., 2022) contains machine translation (MT) tasks for Tibetan and Uyghur. We additionally construct an MRC dataset, **XQuAD-MT**, by translating XQuAD (Artetxe et al., 2020) into Uyghur and Tibetan with Google Translate⁶. All the evaluated datasets only consist of testing sets. See data statistics in Appendix D.2.

Settings For smaller encoder-only models, we adopt **zero-shot transfer**, i.e., finetuning a model with instances in a high-resource language and directly testing it on the low-resource language instances (Artetxe et al., 2020). For WCM-v2, we use its Chinese instances for training. For TibetanQA and XQuAD-MT, we use CMRC 2018 (Cui et al., 2019), a Chinese MRC dataset for

training. For larger generative models, we adopt **in-context learning**, providing exemplars in the prompt (Brown et al., 2020). See the prompts in Appendix D.3.

Compared Models We compare our models with a wide range of multilingual models, including CINO (Yang et al., 2022), mBERT (Devlin et al., 2019), XLM-RoBERTa (Conneau et al., 2019), BLOOM (BigScience et al., 2022), mT5 (Xue et al., 2021) and ByT5 (Xue et al., 2022). These models (potentially) process one or more languages in our study. XLM-RoBERTa is a multilingual pre-trained language model (PLM), whose vocabulary contains 3 Tibetan tokens, 5 Mongolian tokens, and more than 14K tokens in the Arabic script. However, it has only been pretrained in Uyghur. CINO is the largest PLM so far focusing on the minority languages in China, but its pretraining corpora are closed source. BLOOM, mT5, and ByT5 are multilingual LLMs, which can represent tens of languages. Although they are not trained in the four minority languages, their tokenizers can encode unseen languages instead of treating them as unknown tokens. See more details of the supported languages and numbers of parameters in Appendix D.3.

5.3 Results

We show the experiment results of zero-shot transfer in Table 5, with XLM-RoBERTa-large as the baseline. We find both MC²XLMR-large and CINO can outperform vanilla XLM-RoBERTa-large on the classification and MRC tasks. It indicates that pretraining on minority language corpora can effectively enhance the model’s ability to represent these languages. When comparing the two models for Chinese minority languages, we find our MC²XLMR-large can exhibit comparable performance to CINO, which is trained on a closed-source corpus three times larger than MC². This proves the high quality of MC², which can contribute to more transparent and reproducible

⁶<https://translate.google.com/>

Model	WCM-v2	TibetanQA
	Acc. / W.F1	EM / F1
BLOOM-7.1B	9.2 / 8.9	0.0 / 1.3
mT5-xxl	10.0 / 10.4	0.0 / 2.6
ByT5-xxl	10.2 / 10.6	2.5 / 33.1
MC ² Llama-13B (Ours)	36.9 / 37.0	3.8 / 31.5

Table 6: Performance of LLMs under in-context learning. For WCM-v2, we report averaged accuracy and weighted F1 of four languages. For TibetanQA, we use EM and F1 as metrics.

research on minority languages in China.

Table 6 illustrates the results of LLMs under the in-context learning setting. The performance of evaluated multilingual LLMs on WCM-v2 resembles random guessing, with an accuracy of around 10% in the ten-class classification tasks. These models struggle with the Tibetan MRC task. In contrast, the MC²Llama-13B model, pretrained on the MC² corpus, outperforms others by more than +26% accuracy and F1 on WCM-v2. It also achieves 31.5% F1 on TibetanQA. These results show that MC²Llama-13B has a superior ability to handle tasks in the four minority languages of China. Besides smaller models, MC² can also be leveraged for the pretraining of LLMs, thereby effectively enhancing the models’ ability to represent these languages.

The results on Flores-200 and XQuAD-MT are reported in Appendix E. The models trained on MC² also demonstrate superior performance on these tasks.

6 Conclusion

We present MC², an open-source corpus for four minority languages in China. It is also the first corpus focusing on two underrepresented writing systems, i.e., the Kazakh Arabic script and the traditional Mongolian script. To address the severe quality issues in previous low-resource datasets, we adopt a quality-centric strategy to collect MC². We also emphasize the cultural significance of including low-resource writing systems through empirical studies. Our data and models are openly accessible to the community to facilitate research and applications on low-resource languages.

Limitations

Data Sources MC² is mainly composed of web crawls. We have done our best to enrich the domains covered by the MC corpus. However, it is

not feasible to collect data from all websites worldwide. In this work, due to copyright constraints, we are also unable to collect and publicly release a corpus sourced from books in these minority languages. We note that the texts from books can also be used to pretrain language models.

News articles are one of our primary data sources, which are crawled from public websites. The data could potentially contain the biases of these media. Therefore, the data should be used with caution to avoid potential biases and misrepresentations.

Probing Questions It is challenging to collect the kind of probing question in our study, which requires mining the differences between the two communities and formulating them as multi-choice questions. We made great efforts and collected 25 verified probing questions covering a diverse range of topics. As a preliminary study, the size of probing questions is small. We plan to design efficient methods for cultural question sourcing and expand the data scale.

Acknowledgments

This work is supported in part by NSFC (62161160339) and Beijing Science and Technology Program (Z231100007423011). We thank the anonymous reviewers for their valuable suggestions. We thank native speakers of minority languages who contribute to this work. For any correspondence, please contact Yansong Feng.

References

- Julien Abadji, Pedro Ortiz Suarez, Laurent Romary, and Benoît Sagot. 2022. [Towards a cleaner document-oriented multilingual crawled corpus](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4344–4355, Marseille, France. European Language Resources Association.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. [On the cross-lingual transferability of monolingual representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4623–4637, Online. Association for Computational Linguistics.
- Feilong Bao, Guanglai Gao, Xueliang Yan, and Hongxi Wei. 2014. Research on conversion approach between traditional mongolian and cyrillic mongolian. *Computer Engineering and Computation*, 23:206–211.

- BigScience, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. [arXiv preprint arXiv:2211.05100](#).
- A.Z. Broder. 1997. [On the resemblance and containment of documents](#). In [Proceedings, Compression and Complexity of SEQUENCES 1997 \(Cat. No.97TB100171\)](#), pages 21–29.
- 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.
- Samuel Cahyawijaya, Genta Indra Winata, Bryan Wilie, Karissa Vincentio, Xiaohong Li, Adhiguna Kuncoro, Sebastian Ruder, Zhi Yuan Lim, Syafri Bahar, Masayu Khodra, Ayu Purwarianti, and Pascale Fung. 2021. [IndoNLG: Benchmark and resources for evaluating Indonesian natural language generation](#). In [Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing](#), pages 8875–8898, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Bharathi Raja Chakravarthi, Mihael Arcan, and John P McCrae. 2019. Comparison of different orthographies for machine translation of under-resourced dravidian languages. In [2nd Conference on Language, Data and Knowledge \(LDK 2019\)](#). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. [arXiv preprint arXiv:1911.02116](#).
- 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.
- Marta R Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, et al. 2022. No language left behind: Scaling human-centered machine translation. [arXiv preprint arXiv:2207.04672](#).
- Yiming Cui, Ting Liu, Wanxiang Che, Li Xiao, Zhipeng Chen, Wentao Ma, Shijin Wang, and Guoping Hu. 2019. [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 \(EMNLP-IJCNLP\)](#), pages 5883–5889, Hong Kong, China. 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.
- Sumanth Doddapaneni, Rahul Aralikkatte, Gowtham Ramesh, Shreya Goyal, Mitesh M. Khapra, Anoop Kunchukuttan, and Pratyush Kumar. 2023. [Towards leaving no Indic language behind: Building monolingual corpora, benchmark and models for Indic languages](#). In [Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics \(Volume 1: Long Papers\)](#), pages 12402–12426, Toronto, Canada. Association for Computational Linguistics.
- A. Seza Doğruöz, Sunayana Sitaram, Barbara E. Bullock, and Almeida Jacqueline Toribio. 2021. [A survey of code-switching: Linguistic and social perspectives for language technologies](#). In [Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing \(Volume 1: Long Papers\)](#), pages 1654–1666, Online. Association for Computational Linguistics.
- Abteen Ebrahimi and Katharina Kann. 2021. [How to adapt your pretrained multilingual model to 1600 languages](#). In [Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing \(Volume 1: Long Papers\)](#), pages 4555–4567, Online. Association for Computational Linguistics.
- Conghui He, Zhenjiang Jin, Chao Xu, Jiantao Qiu, Bin Wang, Wei Li, Hang Yan, JiaQi Wang, and Dahua Lin. 2023. [Wanjuan: A comprehensive multimodal dataset for advancing english and chinese large models](#). [arXiv preprint arXiv:2308.10755](#).
- Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piñeras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. 2022. [Challenges and strategies in cross-cultural NLP](#). In [Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics \(Volume 1: Long Papers\)](#), pages 6997–7013, Dublin, Ireland. Association for Computational Linguistics.

- Quzhe Huang, Kun Xu, Kun Xu, and Liwei Chen. 2023. Llama2 with chinese continue training. https://huggingface.co/quzhe/llama_chinese_13B.
- Zhengbao Jiang, Antonios Anastasopoulos, Jun Araki, Haibo Ding, and Graham Neubig. 2020. [X-FACTR: Multilingual factual knowledge retrieval from pre-trained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5943–5959, Online. Association for Computational Linguistics.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2016. Bag of tricks for efficient text classification. [arXiv preprint arXiv:1607.01759](#).
- Julia Kreutzer, Isaac Caswell, Lisa Wang, Ahsan Wahab, Daan van Esch, Nasanbayar Ulzii-Orshikh, Allahsera Tapo, Nishant Subramani, Artem Sokolov, Claytone Sikasote, Monang Setyawan, Supheakmongkol Sarin, Sokhar Samb, Benoît Sagot, Clara Rivera, Annette Rios, Isabel Papadimitriou, Salomey Osei, Pedro Ortiz Suarez, Iroro Orife, Kelechi Ogueji, Andre Niyongabo Rubungo, Toan Q. Nguyen, Mathias Müller, André Müller, Shamsuddeen Hassan Muhammad, Nanda Muhammad, Ayanda Mnyakeni, Jamshidbek Mirzakhlov, Tapiwanashe Matangira, Colin Leong, Nze Lawson, Sneha Kudugunta, Yacine Jernite, Mathias Jenny, Orhan Firat, Bonaventure F. P. Dossou, Sakhile Dlamini, Nisansa de Silva, Sakine Çabuk Ballı, Stella Biderman, Alessia Battisti, Ahmed Baruwa, Ankur Bapna, Pallavi Baljekar, Israel Abebe Azime, Ayodele Awokoya, Duygu Ataman, Orevaoghene Ahia, Oghenefego Ahia, Sweta Agrawal, and Mofetoluwa Adeyemi. 2022. [Quality at a glance: An audit of web-crawled multilingual datasets](#). *Transactions of the Association for Computational Linguistics*, 10:50–72.
- Sneha Kudugunta, Isaac Caswell, Biao Zhang, Xavier Garcia, Derrick Xin, Aditya Kusupati, Romi Stella, Ankur Bapna, and Orhan Firat. 2024. Madlad-400: A multilingual and document-level large audited dataset. *Advances in Neural Information Processing Systems*, 36.
- 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. 2022. The bigscience roots corpus: A 1.6 tb composite multilingual dataset. *Advances in Neural Information Processing Systems*, 35:31809–31826.
- Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. 2022. [Deduplicating training data makes language models better](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8424–8445, Dublin, Ireland. Association for Computational Linguistics.
- Bin Li, Yixuan Weng, Bin Sun, and Shutao Li. 2022. A multi-tasking and multi-stage chinese minority pre-trained language model. In *China Conference on Machine Translation*, pages 93–105. Springer.
- Yanchen Liu, William Held, and Diyi Yang. 2023. [DADA: Dialect adaptation via dynamic aggregation of linguistic rules](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13776–13793, Singapore. Association for Computational Linguistics.
- Jo McCormack, Murray Pratt, and Alistair Rolls Alistair Rolls. 2011. [Hexagonal variations: diversity, plurality and reinvention in contemporary France](#), volume 359. Rodopi.
- Benjamin Muller, Antonios Anastasopoulos, Benoît Sagot, and Djamel Seddah. 2021. [When being unseen from mBERT is just the beginning: Handling new languages with multilingual language models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 448–462, Online. Association for Computational Linguistics.
- Preslav Nakov and Hwee Tou Ng. 2009. [Improved statistical machine translation for resource-poor languages using related resource-rich languages](#). In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 1358–1367, Singapore. Association for Computational Linguistics.
- Thuat Nguyen, Chien Van Nguyen, Viet Dac Lai, Hieu Man, Nghia Trung Ngo, Franck Dernoncourt, Ryan A Rossi, and Thien Huu Nguyen. 2023. Culturax: A cleaned, enormous, and multilingual dataset for large language models in 167 languages. [arXiv preprint arXiv:2309.09400](#).
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. The refinedweb dataset for falcon llm: outperforming curated corpora with web data, and web data only. [arXiv preprint arXiv:2306.01116](#).
- Nuo Qun, Xing Li, Xipeng Qiu, and Xuanjing Huang. 2017. End-to-end neural text classification for tibetan. In *Chinese Computational Linguistics and Natural Language Processing Based on Naturally Annotated Big Data: 16th China National Conference, CCL 2017, and 5th International Symposium, NLP-NABD 2017, Nanjing, China, October 13-15, 2017, Proceedings 5*, pages 472–480. Springer.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.

- Surangika Ranathunga, Nisansa De Silva, Velayuthan Menan, Aloka Fernando, and Charitha Rathnayake. 2024. [Quality does matter: A detailed look at the quality and utility of web-mined parallel corpora](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 860–880, St. Julian's, Malta. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang, Suraj Srivats, Soroush Vosoughi, Hyung Won Chung, Yi Tay, Sebastian Ruder, Denny Zhou, Dipanjan Das, and Jason Wei. 2022a. [Language models are multilingual chain-of-thought reasoners](#). Preprint, arXiv:2210.03057.
- Hanru Shi, Sisi Liu, Xinhe Yu, Wugede Bao, Yuan Sun, and Xiaobing Zhao. 2022b. [Milmo: Minority multilingual pre-trained language model](#). arXiv preprint arXiv:2212.01779.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. [Megatron-lm: Training multi-billion parameter language models using model parallelism](#). arXiv preprint arXiv:1909.08053.
- Simeng Sun, Angela Fan, James Cross, Vishrav Chaudhary, Chau Tran, Philipp Koehn, and Francisco Guzmán. 2022. [Alternative input signals ease transfer in multilingual machine translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5291–5305, Dublin, Ireland. Association for Computational Linguistics.
- Yuan Sun, Sisi Liu, Chaofan Chen, Zhengcuo Dan, and Xiaobing Zhao. 2021. [Teaching machines to read and comprehend tibetan text](#). *Journal of Computer and Communications*, 9(09):143–152.
- Daniela Teodorescu, Josie Matalski, Delaney Lothian, Denilson Barbosa, and Carrie Demmans Epp. 2022. [Cree corpus: A collection of nêhiyawêwin resources](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6354–6364, Dublin, Ireland. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shrusti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). arXiv preprint arXiv:2307.09288.
- Genta Winata, Alham Fikri Aji, Zheng Xin Yong, and Tamar Solorio. 2023. [The decades progress on code-switching research in NLP: A systematic survey on trends and challenges](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2936–2978, Toronto, Canada. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2022. [ByT5: Towards a token-free future with pre-trained byte-to-byte models](#). *Transactions of the Association for Computational Linguistics*, 10:291–306.
- 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.
- Ziqing Yang, Zihang Xu, Yiming Cui, Baoxin Wang, Min Lin, Dayong Wu, and Zhigang Chen. 2022. [CINO: A Chinese minority pre-trained language model](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3937–3949, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Zheng Xin Yong, Hailey Schoelkopf, Niklas Muenighoff, Alham Fikri Aji, David Ifeoluwa Adelani, Khalid Almubarak, M Saiful Bari, Lintang Sutawika, Jungo Kasai, Ahmed Baruwa, Genta Winata, Stella Biderman, Edward Raff, Dragomir Radev, and Vasilina Nikoulina. 2023. [BLOOM+1: Adding language support to BLOOM for zero-shot prompting](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11682–11703, Toronto, Canada. Association for Computational Linguistics.
- Marcos Zampieri, Preslav Nakov, and Yves Scherrer. 2020. [Natural language processing for similar languages, varieties, and dialects: A survey](#). *Natural Language Engineering*, 26(6):595–612.
- Chen Zhang, Xiao Liu, Jiaheng Lin, and Yansong Feng. 2024. [Teaching large language models an unseen language on the fly](#). arXiv preprint arXiv:2402.19167.

Sheng Zhang, Xin Zhang, Weiming Zhang, and Anders Søgaard. 2021. [Sociolectal analysis of pre-trained language models](#). In [Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing](#), pages 4581–4588, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Caleb Ziems, Jiaao Chen, Camille Harris, Jessica Anderson, and Diyi Yang. 2022. [VALUE: Understanding dialect disparity in NLU](#). In [Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics \(Volume 1: Long Papers\)](#), pages 3701–3720, Dublin, Ireland. Association for Computational Linguistics.

A Corpus Details

A.1 Quality Issues in Previous Corpora

We provide several cases of quality issues in previous corpora.

Language Misidentification In the Uyghur split of CulturaX, the website `kazakh.people.com.cn` ranks third in terms of document count, comprising 9% of the total documents in the dataset. However, this website is a Kazakh-language news website. These documents are mislabeled as Uyghur because the website uses the Kazakh Arabic script, which is similar to the Uyghur Arabic script.

Insufficient Data Cleaning In Figure 5, we show an example of insufficient data cleaning from the Tibetan corpus in CulturaX. The highlighted parts are irrelevant texts on a web page, such as headers and footers.

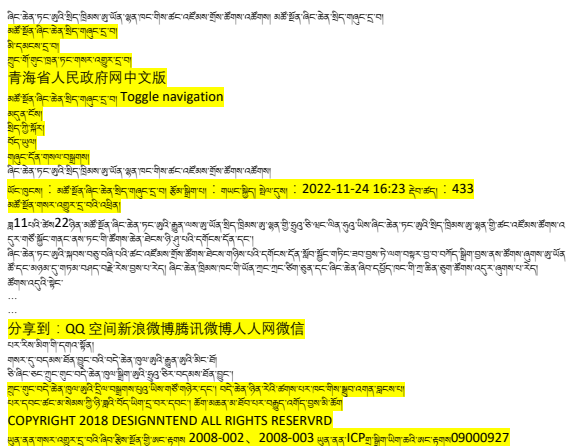


Figure 5: An example of insufficient data cleaning from the Tibetan corpus in CulturaX. The highlighted parts are irrelevant texts on a web page, such as headers and footers.

A.2 Deduplication and Filtering

We describe in detail the deduplication and filtering process of MC².

Deduplication We first discard the web pages from different sources with the same URLs. We then adopt both exact and fuzzy deduplication. For exact deduplication, we use the SHA-256 hash function to calculate the hash value of each web page’s content and keep only one page for the pages with the same hash value. For fuzzy deduplication, we remove similar web pages by applying Min-Hash (Broder, 1997): for each page, we compute the minhash values and measure its approximate

Subset	%
Tibetan	2.5%
Uyghur	9.3%
Kazakh	0.9%
Mongolian	0.7%

Table 7: Percentage of documents containing Chinese characters for the four subsets in MC².

similarity with other pages, removing pairs whose minhash values are the same in at least one bucket. We use the same parameters as Lee et al. (2022): $n = 5$ (5-grams), $b = 450$, and $r = 20$.

Filtering We use the following rules:

- **Repetition:** We remove the documents that have a high ratio of paragraph repetition or n-gram repetition.
- **Document Length:** We remove the documents whose lengths are below the threshold.
- **Unexpected Character:** We remove the document whose ratios of the characters in the target language are below the threshold.

Note that the hyperparameters in the heuristic for filtering English corpora are not necessarily applicable to low-resource languages due to the difference in encodings and linguistic characteristics. One needs to decide the hyperparameters for each language respectively.

A.3 Code-Switching in MC²

Code-switching is an important and interesting phenomenon in multilingual NLP. Here we share some analyses and findings of the code-switching phenomenon present in our corpus. We mainly focus on the code-switching between Chinese and the four minority languages in China.

First, we calculate the proportion of documents containing Chinese characters for the four languages in MC². The results are shown in Table 7. 3.3% of the documents in MC² contain Chinese characters, indicating the commonness of the code-switching phenomenon in our corpus. Uyghur exhibits the highest rate of code-switching. We find that the Uyghur documents in our corpus are more likely to attach Chinese translations after entity mentions.

Second, we analyze the types of code-switching present in our corpus. Inspired by previous works (Doğruöz et al., 2021; Winata et al., 2023),

Type	bo	ug	kk	mn
Intra-sentential	6%	14%	38%	36%
Inter-sentential	46%	0%	28%	24%
Attached translations	48%	86%	34%	40%

Table 8: Percentage of documents containing Chinese characters for the four subsets in MC².

we classify documents containing Chinese characters into three types:

- Intra-sentential code-switching: switching that occurs in a sentence where a word was switched to another language.
- Inter-sentential code-switching: switching that occurs in two different sentences.
- Attached translations of entities: In many cases, authors of articles often append Chinese translations after entity mentions. Strictly speaking, this doesn't fall under code-switching, but this situation is quite common.

We sample 50 documents containing Chinese characters and manually classify them into three types. The proportion of each type is shown in Table 8. The most common phenomenon is attaching Chinese translations to the entities mentioned in the documents. The common entities under these circumstances include scientific and cultural terms. Regarding intra-sentential code-switching, the switched phrases are often entities of Chinese origin, such as Han Chinese names and literary works written in Chinese. These Chinese terms currently lack official translations in these languages, so they are directly borrowed in their original form. Many instances of inter-sentential code-switching are sentences in different languages, where the sentences before and after each other are translations of one another. Based on this property, we can mine parallel sentences from the corpus, which we leave as future work.

B Cultural Probing

We provide a detailed explanation of the probing cases in Table 4. For Query 1, May 1st is a public holiday in both China and Kazakhstan. However, the people in the two countries celebrate different festivals. In China, people celebrate International Workers' Day on May 1st, to commemorate the struggles of laborers and the working class for the eight-hour workday. However, according to

the Kazakhstan law signed on October 18, 1995, May 1st was renamed People's Unity Day, and the Soviet-era Labor Day was formally canceled.

For Query 2, in mainland China, the official currency is Renminbi. Meanwhile, in Kazakhstan, the official currency is Kazakhstani Tenge.

For Query 3, Sayram Lake is an endorheic freshwater lake in the northern Tianshan Mountains, near to Ili Kazakh Autonomous Prefecture in China. Its name, *Sayram*, is believed to originally derive from Kazakh, meaning *blessing*. Sayram Lake is famous in China as *the last drop of the Atlantic tears*. Lake Balkhash is also famous in Kazakhstan since it is the largest lake in Kazakhstan. Around 20% of Kazakhstani people are living in the drainage basin of Lake Balkhash, including the residents of Almaty, the largest city in Kazakhstan.

C Mongolian Transliteration

We use an open-source tool to transliterate Cyrillic Mongolian into traditional Mongolian⁷. It is trained on 80K sentence pairs crawled from a closed-source transliteration system. When transliterating a Cyrillic Mongolian corpus, we find that the transliteration quality deteriorates with the increase in the input length. So we split each document into sentences and do the transliteration at a sentence level to improve the quality. However, we find that the transliteration results are imperfect, especially in the case of long sentences. In Figure 6, the transliteration tool produces a bad result for a long input sentence, repeating similar words at the end of the output.

D Experiment Details

D.1 Model Training

Hyperparameters We list in Table 9 the hyperparameters for the continual pretraining of MC²XLMR-large and MC²Llama-13B.

Training Frameworks For MC²XLMR-large, we use Transformers (Wolf et al., 2020) for training. For MC²Llama-13B, we use Megatron (Shoeybi et al., 2019) for training.

Computing Infrastructure We continually pre-train MC²XLMR-large on one A100 GPU, and an epoch takes 17 hours. We continually pre-train MC²Llama-13B on eight A100 GPUs, and an epoch takes 20 hours.

⁷<https://github.com/tugstugi/mongolian-nlp/tree/master/bichig2cyrillic>

models sometimes fail to output a label in the label list. In these cases, we assign random labels to these samples.

E Additional Results

E.1 WCM-v2

In Table 13, we report the performance in different languages on WCM-v2.

E.2 Flores-200

Considering that encoder-only models cannot perform the MT task, we use generative models for experiments. Since the models have limited abilities to generate texts in low-resource languages, we only conduct experiments to translate the low-resource languages into Chinese.

The results are shown in Table 14. As they are not finetuned for the translation task, the four models generally perform poorly on this task merely through in-context learning. However, our model achieves non-zero scores, demonstrating its preliminary ability to understand these languages, thanks to the training on MC². We leave the supervised training setting of MT for future work.

E.3 XQuAD-MT

In Table 15 we report the performance on the full set and EXACT subset of XQuAD-MT.

Model	Params	Supported Languages			
		Tibetan	Uyghur	Kazakh (Arabic)	Mongolian (Traditional)
mBERT-base	110M	✗	✗	✗	✗
XLNet-RoBERTa-large	560M	✗	✓	✗	✗
CINO-large-v2	442M	✓	✓	✓	✓
BLOOM-7.1B	7.1B	✗	✗	✗	✗
mT5-xxl	13B	✗	✗	✗	✗
ByT5-xxl	13B	✗	✗	✗	✗
MC ² XLNet-large (Ours)	572M	✓	✓	✓	✓
MC ² Llama-13B (Ours)	13B	✓	✓	✓	✓

Table 11: The supported languages of the models in our experiments and their parameters.

WCM-v2 (Classification)

请仿照示例，将下面的{藏语/维吾尔语/哈萨克语/蒙古语}文本分类为艺术、地理、历史、自然、自然科学、人物、技术、教育、经济、健康中的一个类别。

Follow the example and classify the following {Tibetan/Uyghur/Kazakh/Mongolian} passage into one of the categories of Art, Geography, History, Nature, Natural Science, People, Technology, Education, Economy, and Health.

示例{1/2/3} *Example {1/2/3}*

文本: {Passage written in Tibetan} *Passage: {Passage written in Tibetan}*

类别: {Class} *Class: {Class}*

请仿照以上示例，将下面的{藏语/维吾尔语/哈萨克语/蒙古语}文本分类为艺术、地理、历史、自然、自然科学、人物、技术、教育、经济、健康中的一个类别。

Follow the example above to classify the following {Tibetan/Uyghur/Kazakh/Mongolian} passage into one of the categories of Art, Geography, History, Nature, Natural Science, People, Technology, Education, Economy, and Health.

TibetanQA (Question Answering)

阅读文章，从文章中抽取内容回答问题。文章、问题和答案都应是藏文。接下来是两个样例。

Read the passage and answer the questions by extracting content from the passage. Passages, questions, and answers should be written in Tibetan. Here are two examples.

样例{1/2} *Example {1/2}*

文章: {Context written in Tibetan} *Passage: {Context written in Tibetan}*

问题: {Question written in Tibetan} *Question: {Question written in Tibetan}*

答案: {Answer written in Tibetan} *Answer: {Answer written in Tibetan}*

Table 12: Prompt templates used in our experiments. The *text in italics* are the English translations of the Chinese instructions.

Model	Open-source Corpus	Classification (Accuracy / Weighted F1)				
		bo	ug	kk	mn	All
BLOOM-7.1B	Yes	11.8 / 9.5	5.7 / 8.9	8.4 / 7.4	9.9 / 13.0	9.2 / 8.9
mT5-xxl	Yes	9.2 / 11.7	13.3 / 20.1	10.6 / 11.3	10.1 / 13.6	10.0 / 10.4
ByT5-xxl	Yes	10.7 / 12.7	8.6 / 14.7	10.3 / 10.9	9.9 / 12.9	10.2 / 10.6
MC ² Llama-13B (Ours)	Yes	30.7 / 33.7	39.0 / 52.5	34.0 / 32.4	45.4 / 51.9	36.9 / 37.0

Table 13: Performance of different models on WCM-v2 under in-context learning. The best scores are made **bold**. We use accuracy and weighted F1 as metrics.

Model	Tibetan	Uyghur
BLOOM-7.1B	0.4 / 1.9	0.9 / 3.4
mT5-xxl	0.0 / 0.3	0.0 / 0.4
MC ² Llama-13B (Ours)	2.5 / 4.4	6.5 / 8.2

Table 14: Performance of different models on Flores-200 under in-context learning. The best scores are made **bold**. The scores in each cell are BLEU and chrF.

Model	Uyghur		Kazakh	
	Full	EXACT	Full	EXACT
mBERT-base	0.2 / 3.8	0.8 / 6.2	0.3 / 6.0	1.0 / 9.6
XLM-RoBERTa-large	2.4 / 9.5	10.9 / 20.9	2.2 / 13.5	<u>6.5</u> / 24.0
CINO-large-v2	1.8 / 9.8	8.6 / <u>20.7</u>	2.7 / <u>13.9</u>	7.7 / <u>24.9</u>
MC ² XLMR-large (Ours)	<u>2.0</u> / <u>9.6</u>	<u>9.3</u> / 19.5	<u>2.3</u> / 16.4	<u>6.5</u> / 27.0

Table 15: Performance of different models on XQuAD-MT under zero-shot transfer. The best scores are made **bold**, with the second underlined. The scores in each cell are EM and F1.