

# Do Large Language Models Have an English “Accent”?

## Evaluating and Improving the Naturalness of Multilingual LLMs

Yanzhu Guo<sup>2,3\*</sup>, Simone Conia<sup>4</sup>, Zelin Zhou<sup>1</sup>, Min Li<sup>1</sup>, Saloni Potdar<sup>1</sup>, Henry Xiao<sup>1</sup>  
<sup>1</sup>Apple <sup>2</sup>Inria Paris <sup>3</sup>École Polytechnique <sup>4</sup>Sapienza University of Rome

### Abstract

Current Large Language Models (LLMs) are predominantly designed with English as the primary language, and even the few that are multilingual tend to exhibit strong English-centric biases. Much like speakers who might produce awkward expressions when learning a second language, LLMs often generate unnatural outputs in non-English languages, reflecting English-centric patterns in both vocabulary and grammar. Despite the importance of this issue, the naturalness of multilingual LLM outputs has received limited attention. In this paper, we address this gap by introducing novel automatic corpus-level metrics to assess the lexical and syntactic naturalness of LLM outputs in a multilingual context. Using our new metrics, we evaluate state-of-the-art LLMs on a curated benchmark in French and Chinese<sup>1</sup>, revealing a tendency towards English-influenced patterns. To mitigate this issue, we also propose a simple and effective alignment method to improve the naturalness of an LLM in a target language and domain, achieving consistent improvements in naturalness without compromising the performance on general-purpose benchmarks. Our work highlights the importance of developing multilingual metrics, resources and methods for the new wave of multilingual LLMs.

### 1 Introduction

LLMs are becoming an integral part of society with a visible impact on the general public (Heikkilä, 2023; Ziegler and Donkers, 2024; Peterson, 2024). However, there exists a significant disparity in how different languages are represented in LLMs, as popular models are primarily designed with English in mind. While this is a well-known issue and more multilingual LLMs are being released, most of them are still English-dominated (Wendler

et al., 2024). A prominent example is the Llama 3.1 series of models (Dubey et al., 2024), which are claimed to be state-of-the-art multilingual LLMs: these models are trained on 15T tokens, yet only 8% of the training data is declared to be non-English.

We could make an analogy between these multilingual LLMs and native English speakers who are trying to acquire a new language (Groves et al., 2018; DeVore and Kyle, 2023). Their language notions are built in an English-centric system, and they inevitably bring traces of English habits into other languages when transferring their notions (Papadimitriou et al., 2023; Wendler et al., 2024). Moreover, because of the lack of data in non-English languages, multilingual LLMs are often exposed – either during pre-training or post-training (Yang et al., 2024; Abdin et al., 2024) – to texts translated from English. Both human and machine translated language are known to suffer from translationese artifacts, which set them apart from native content (Bizzoni et al., 2020; Luo et al., 2024). LLMs trained on such data are susceptible to suffer from the same translationese problems.

In fact, even LLM-generated texts in English are known to exhibit distributional differences from human-written texts (Guo et al., 2023; Liang et al., 2024a,b; Shumailov et al., 2024). Given the predominance of English data, this effect is likely more pronounced in non-English outputs. However, current evaluations of multilingual LLMs still focus on their task-solving capabilities (Zheng et al., 2023; Feng et al., 2023; Zhang et al., 2024; Hendrycks et al., 2021), overlooking the aspect of naturalness. As LLMs increasingly influence various aspects of our lives, their tendency to produce less natural outputs in lower-resource languages in favor of English expressions could amplify the unfairness for the communities that speak these languages. Therefore, it is crucial to evaluate and improve the naturalness of multilingual LLMs to foster fair language representation.

\*Work done during internship at Apple.

<sup>1</sup>Chinese in this paper refers to Mandarin written in Simplified Chinese characters.

In this paper, we take two steps towards this goal. Our first step is the introduction of a novel set of metrics to evaluate the naturalness of LLM outputs at a corpus level by comparing the lexical and syntactic distributions of LLM outputs with human written texts. We use these metrics, together with our new topically-aligned cross-lingual dataset, to benchmark and analyze the naturalness of state-of-the-art multilingual LLMs in English, French and Chinese. Our second step is the introduction of a simple and effective approach to enhancing the naturalness of multilingual LLMs in a target language. Using Direct Preference Optimization (DPO) (Rafailov et al., 2023), we leverage a new preference dataset that contrasts human-written responses with synthetically-manipulated ones. Experimental results show that our method consistently improves naturalness of an LLM in Chinese without sacrificing its capabilities on general-purpose benchmarks.

In summary, our contributions are threefold:

1. We develop new metrics to evaluate the lexical and syntactic naturalness of LLM outputs in a multilingual setting;
2. We create a benchmark for cross-lingual evaluation of LLM naturalness and draw insights from the benchmark results on important factors that could impact LLM naturalness;
3. We propose an alignment approach for improving the naturalness of existing LLMs with promising results across models and domains.

We hope our investigation will encourage further research on the limitations of multilingual LLMs beyond their scores on task-solving benchmarks.

## 2 Related Work

Although the evaluation of LLMs has received significant research interest in recent years, the greater part of this body of work has focused on aspects such as helpfulness (Zheng et al., 2023), factual accuracy (Feng et al., 2023), safety (Zhang et al., 2024), fairness (Chalkidis et al., 2022), and task-specific performance (Hendrycks et al., 2021), leaving the naturalness of LLM outputs under-investigated. Naturalness is commonly used as an evaluation criterion in machine translation, but it has mostly relied on either human ratings (Chen et al., 2024) or trained classifiers (Liu et al., 2021). To the best of our knowledge, our work is the first

to systematically investigate linguistic naturalness in multilingual LLM generations outside of the machine translation context. Nonetheless, several adjacent research areas are highly relevant to our focus, including translationese detection, linguistic diversity evaluation, and multilingual language model analysis.

**Translationese detection**, a well-established task in machine translation, aims to determine whether a text is originally written in the target language or translated from another language (Volansky et al., 2015; Wintner, 2016). For instance, Freitag et al. (2022) employ a range of linguistic features – including type-token ratio, lexical density, answer lengths, dependency tree height, constituency tree height, perplexity, etc. – to train a classifier to distinguish between machine-translated and naturally occurring sentences. However, these classifiers are prone to overfit on specific training data. In an effort to enhance the naturalness of machine translation outputs, Freitag et al. (2022) also tagged parallel training data based on target-side naturalness, contrasting models trained on natural versus translated text. Our approach builds on the concept of contrasting natural versus unnatural texts but avoids reliance on pre-trained classifiers. Instead, we use automatically manipulated unnatural texts and preference learning. We also extend the analysis beyond machine translation to general multilingual text generation.

**Linguistic diversity evaluation** (Tevet and Berant, 2021) is another area closely related to naturalness, as a key marker of unnaturalness in synthetic text is reduced linguistic diversity (Vanmassenhove et al., 2021; Guo et al., 2024). Our work draws inspiration from the way these diversity features are computed. Past work compares the diversity of generated texts to human texts and consider texts to be natural if they approach the human level of diversity (Freitag et al., 2022). However, such an approach primarily assesses the dispersion of the distribution while overlooking more holistic comparisons of linguistic features. To address this, we introduce metrics that directly compare vocabulary and syntactic distributions between human and machine-generated texts.

**Multilingual analysis** of LLMs has recently shown that models trained on unbalanced, English-heavy corpora often rely on English as an internal pivot language. For instance, Wendler et al. (2024) demonstrate that the concept space of LLaMA-2 is

more closely aligned with English than with other input languages. More empirically, Papadimitriou et al. (2023) show that multilingual BERT exhibits a bias toward English-like grammatical structures. Despite these findings, there has been no systematic study on how this English-centric tendency affects the linguistic naturalness of multilingual LLM outputs, particularly in open-ended downstream tasks.

### 3 Evaluation Metrics for Naturalness

In this section, we introduce a new set of evaluation metrics designed to assess the naturalness of multilingual text generation at the corpus level. While our approach also requires a reference set of natively written texts, it differs from widely used reference-based metrics such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), and BERTScore (Zhang\* et al., 2020). These sample-level reference metrics often struggle to account for human label variability and uncertainty (Plank, 2022; Giulianelli et al., 2023), particularly in open-ended tasks with multiple valid generations. Additionally, while a single text sample with certain choices of vocabulary or grammatical structure might seem natural, their repeated occurrence across many generations would raise a red flag (Guo et al., 2023). Detecting unnaturalness in single instances can be difficult, but statistical patterns emerge more clearly at the corpus level, serving as stronger indicators (Liang et al., 2024a,b). Our metrics leverage these distributional patterns, offering a broader and more robust perspective on language use across large text collections.

Having highlighted the advantage of a distribution-level perspective, we propose a new definition for language model naturalness. Past studies have defined the naturalness of a single piece of text by asking “*could it have been produced by a native speaker?*” (Novikova et al., 2016; Groves et al., 2018; Liu et al., 2021). We adapt this definition to the corpus level and define the naturalness of a language model by asking “*could the set of texts generated by this language model have been produced by a group of native speakers?*”

Our metrics are inspired by divergence measures such as MAUVE (Pillutla et al., 2021; Pimentel et al., 2023), which quantify the information-theoretic divergence between the probability distributions of a language generator and a true natural language distribution. However, our approach dif-

fers by not relying on another language model to embed the generated texts. This reduces the risk of introducing intrinsic biases from the chosen embedding model, a crucial consideration in multilingual settings, where such the chosen embedding models are often English-dominated themselves. Our method is also more transparent and interpretable, as it clearly distinguish between two key aspects of linguistic naturalness: syntactic and lexical naturalness. However, our metrics focus exclusively on the linguistic form of the text and do not address semantic aspects. For analyzing semantics, the use of external embedding models may be inevitable.

In the following, we introduce the methodology for evaluating lexical and syntactic naturalness. The implementations of the metrics are described in Appendix C.

#### 3.1 Lexical Naturalness

We propose to measure the lexical naturalness of an LLM by comparing the vocabulary distribution of its generated text with that of human-written text. More specifically, we put forward a lexical naturalness metric based on computing the **Jensen-Shannon Divergence (JSD)** between the lexical distributions of LLM-generated and human-written text from the same prompts. The JSD provides a symmetric and bounded measure of difference between two distributions without them necessarily sharing the same support, i.e., the tokens in the vocabulary of the LLM in our case.<sup>2</sup> Given the vocabulary distributions  $P$  and  $Q$  corresponding to human and model outputs, respectively, the JSD is calculated as follows:

$$\text{JSD}(P||Q) = \frac{1}{2} (D_{\text{KL}}(P||M) + D_{\text{KL}}(Q||M)),$$

where  $M = \frac{1}{2}(P + Q)$  is the midpoint distribution, and  $D_{\text{KL}}$  is the Kullback-Leibler divergence. By assessing the divergence, we can quantify how closely the model’s vocabulary distribution aligns with human language, where lower values indicate greater similarity and higher lexical naturalness. We note that our approach to directly compare the two vocabulary distributions implicitly captures the lexical statistical tendencies of past work that used type token-ratio and rank-frequency coefficient (Meister and Cotterell, 2021).

<sup>2</sup>We process vocabularies at the word level instead of sub-word/token level to represent more meaningful lexical units.

### 3.2 Syntactic Naturalness

To measure the syntactic naturalness of LLMs generations we leverage the **Universal Dependencies (UD)** grammar framework (Nivre et al., 2020). UD provides well-defined, theoretically-grounded linguistic structures across different languages, making it suitable for cross-lingual comparisons. Our proposed metric is based on representing each sentence as a dependency tree, where nodes correspond to words and edges specify the dependency relations between them. Additionally, each word is annotated with its corresponding Part-of-Speech (POS) tag as the node label. At a high-level, our metric computes the structural similarity of all pairs of sentences in a corpus, clustering the ones who share similar syntax, allowing us to determine if there is a distributional difference between two groups (LLM-generated text and human-written text, in our case).

To compute the structural similarity between pairs of dependency trees, we propose to use the **Weisfeiler-Lehman (WL)** graph kernel (Shervashidze et al., 2011). The WL kernel iteratively refines node labels based on the labels of neighboring nodes, thereby constructing a hierarchical encoding of the graph structure. More formally, let  $T_1$  and  $T_2$  be two dependency trees with respective vertex, i.e., word, sets  $V_1$  and  $V_2$ . The WL kernel,  $K_{\text{WL}}(T_1, T_2)$ , measures the similarity between  $T_1$  and  $T_2$  by comparing their subtree structures. Initially, each node in the trees is assigned a label based on its original POS tag. Then, at each iteration  $h$ , we aggregate the labels of the node’s neighbors into a multiset, which is then hashed into a new unique label. This process continues for  $H$  iterations, and the kernel value is computed as the number of matching node labels across all iterations:

$$K_{\text{WL}}(T_1, T_2) = \sum_{h=0}^H \sum_{(v_1, v_2) \in (V_1, V_2)} \delta(\ell_h(v_1), \ell_h(v_2)),$$

where  $\ell_h(v)$  is the label of node  $v$  at iteration  $h$ , and  $\delta(\cdot, \cdot)$  is the Kronecker delta function. For our experiments, we fixed the number of iterations  $H$  to 2, as this proved to be the most effective hyperparameter across various discriminative tasks for graphs (Shervashidze et al., 2011).

Given a set of human-generated sentences  $\{s_i^h\}_{i=1}^{N_h}$  and model-generated sentences  $\{s_j^m\}_{j=1}^{N_m}$ , we compute the WL kernel similarity between all

pairs of dependency trees from these sets. The resulting kernel matrix  $\mathbf{K} \in \mathbb{R}^{N_h \times N_m}$  has elements  $K_{ij} = K_{\text{WL}}(T_i^h, T_j^m)$ , where  $T_i^h$  and  $T_j^m$  represent the dependency trees corresponding to sentences  $s_i^h$  and  $s_j^m$ , respectively. This kernel matrix  $\mathbf{K}$  captures structural similarities between the dependency trees of human-written and LLM-generated sentences.

Once the kernel matrix is obtained, we use the **Maximum Mean Discrepancy (MMD)** (Gretton et al., 2012) to compare the distributions of human-generated and model-generated sentences. In particular, the MMD<sup>2</sup> between the two sets of sentences is computed as:

$$\frac{1}{N_h^2} \sum_{i, i'} K_{ii'} + \frac{1}{N_m^2} \sum_{j, j'} K_{jj'} - \frac{2}{N_h N_m} \sum_{i, j} K_{ij},$$

where  $K_{ii'}$  and  $K_{jj'}$  represent the similarities within the human and model-generated sentence sets, and  $K_{ij}$  represents the cross-similarities between the two sets. The resulting value provides a measure of syntactic divergence between the human-generated and model-generated sentences, with a lower value indicating greater similarity.

To conclude, our proposed syntactic naturalness metric quantifies the syntactic divergence between human and model-generated sentences by examining the distribution of their dependency tree structures. This approach considers both the dependency relationships and the hierarchical arrangement of substructures at multiple levels. Each step in our approach (e.g., the POS tagger, WL kernel, and MMD) has been externally validated for accuracy and effectiveness.

## 4 Cross-lingual Analysis of Naturalness

After introducing our new metrics, we proceed with a cross-lingual analysis of the naturalness of multilingual LLMs. This process involves curating a new dataset, selecting the models, and analyzing the benchmark results.

### 4.1 Dataset and Evaluation

For our analysis, we require a corpus that is typically aligned across languages while preserving ground-truth naturalness. This means that the texts in each language must be natively written, not translated – whether by humans or models. As a result, the parallel corpora typically used in machine translation research are not suitable. Instead, we construct a new dataset that satisfies our criteria start-

Model Size (# parameters)		Human	Qwen1.5	Qwen2	Mistral-v0.3	Mistral-Nemo	Llama-3	Llama-3.1
		—	7B	7B	7B	12B	8B	8B
<b>English</b>	Lexical Divergence	<b>23.07</b>	<b>30.36</b>	25.31	<b>23.30</b>	25.12	29.00	26.79
	Syntactic Divergence	<b>3.53</b>	<b>22.19</b>	13.67	<b>13.56</b>	14.77	17.72	16.80
<b>Chinese</b>	Lexical Divergence	<b>25.91</b>	<b>41.00</b>	37.08	39.02	34.78	36.88	<b>33.29</b>
	Syntactic Divergence	<b>2.93</b>	<b>23.33</b>	20.66	17.29	12.84	15.45	<b>10.32</b>
<b>French</b>	Lexical Divergence	<b>24.25</b>	<b>38.35</b>	31.18	<b>28.73</b>	31.34	32.22	31.52
	Syntactic Divergence	<b>3.22</b>	<b>24.21</b>	12.10	12.72	14.72	17.88	<b>11.27</b>

Table 1: Benchmark results for the lexical and syntactic naturalness of multilingual LLMs. All divergence values are presented as percentages and lower values indicate better naturalness. The best results for each language are highlighted in blue, while the worst are highlighted in orange.

ing from Wikipedia<sup>3</sup>, which offers a wealth of texts that are frequently edited by native speakers and that we can topically-align across languages. Importantly, to minimize cultural bias in our dataset, we select the most-viewed Wikipedia entries across English, Chinese and French.<sup>4</sup> Our data curation process results in 3,722 Wikipedia entries, each accompanied by descriptions in the three target languages. We discuss the preprocessing of this dataset in Appendix A.

For each entry in the dataset, we instruct the models to complete a straightforward task: generating a description for the given entry in each of the three languages. We select this task to avoid overly constraining the content of the language model outputs, allowing them to generate in a more natural and organic manner. We stress that our focus is on the overall distribution of vocabulary and grammatical structures rather than on individual outputs. The prompts and generation settings used in this task are provided in Appendix B.

## 4.2 Multilingual LLMs

We experiment with three families of open-weight LLMs: Llama (Dubey et al., 2024), Qwen (Yang et al., 2024), and Mistral (Jiang et al., 2023). These models are selected for their state-of-the-art performance on diverse English and multilingual benchmarks. Additionally, they are developed by teams from regions where English, Chinese, and French are the official languages, respectively. While the exact composition of the training data for each model is not publicly available, we speculate that all are predominantly English-centric, though we

expect Qwen to include more Chinese data and Mistral to include more French data than the others. We benchmark the two most recent versions from each model family to track changes in naturalness performance over time. For all models, we test their moderately scaled versions, using open-source implementations from the Transformers library (Wolf et al., 2020). We focus on the instruction-tuned versions of the models, as these are more commonly used in real-world applications compared to the base versions.

## 4.3 Results and Analysis

Table 1 presents the naturalness performance of all models based on our proposed metrics (introduced in Section 3). As expected, the results show a consistent improvement in naturalness across newer versions of LLMs compared to their earlier counterparts (e.g., from Qwen-1.5 to Qwen-2). However, our metrics also reveal a persistent gap between the naturalness of human-written and LLM-generated texts, especially in non-English outputs, supporting our initial hypothesis.

Note that the human reference value is computed by randomly selecting non-overlapping subsets of human-written texts and measuring the divergence between them. These values are not obtained under the exact same conditions as the human-model divergences, which may involve overlapping prompts. This stricter setup makes the human reference value we report an upper bound on the human divergence baseline. As we show below, even this upper bound remains substantially lower than the human-model divergences.

**Gap in lexical naturalness.** In terms of lexical divergence, English outputs – especially from Mistral-v0.3 – approach human reference values, while larger gaps persist in Chinese and French,

<sup>3</sup>Wikipedia is available under the CC BY-SA 4.0 license.

<sup>4</sup>We use the Wikipedia API to retrieve the top-20 most-viewed entries in each language daily from June 2022 to June 2025. We then take the per-language union of the entries and retain their the cross-language intersection.

POS Pattern	$n$ -grams from Llama-3.1		Frequency			Explanation
	English	Chinese	Native English	Native Chinese	Llama-3.1 Chinese	
(ADJ, CONJ, ADJ)	(blue, and, white) (incoming, and, outgoing) (eastern, and, north)	(真诚, 和, 直接) (正常, 和, 合法) (血腥, 和, 黑暗)	37	4	22	English requires conjunctions when listing adjectives, while Chinese often omits them.
(PRON, AUX, VERB)	(They, were, married) (he, was, named) (it, was, built)	(他, 被, 释放) (她, 被, 称为) (他, 被, 任命)	175	10	56	English uses passive constructions with auxiliary verbs more frequently than Chinese.
(ADP, DET, NOUN, ADP)	(During, the, boom, of) (at, the, end, of) (At, the, height, of)	(在, 多部, 电影, 中) (在, 此, 职位, 上) (在, 此次, 比赛, 中)	300	3	22	English depends heavily on prepositions and determiners to structure sentences, while Chinese tends to simplify these relationships.

Table 2: Extracted syntactic patterns that show the influence of English structures on how Llama-3.1 generates Chinese text. All frequencies are based on 40,000  $n$ -grams of the same  $n$ . The Chinese and English outputs provided are neither semantically aligned nor translations.

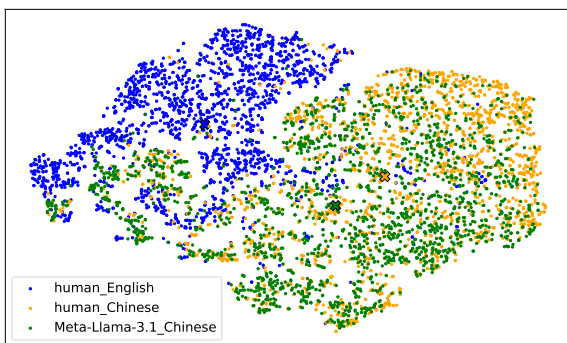


Figure 1: T-SNE visualization of the syntactic structures generated by Llama-3.1 in Chinese, compared to human-written Chinese and human-written English.

with Chinese showing the most significant difference. This indicates that the models are lexically less natural in languages other than English, and the naturalness gap is more pronounced for languages that are typologically more distant from English. The syntactic divergence values across different languages are not directly comparable. This is due to the fact that the UD grammar parses languages at varying levels of granularity, as evidenced by the lower human divergence value in Chinese compared to English or French.

**English “accent” in syntactic structures.** When examining syntactic divergence, all models in all languages show significant differences from the human reference values. Since dependency trees, when parsed using the UD grammar, share a common space across languages, we can perform cross-lingual comparisons of syntactic structures. In Figure 1, we show Llama-3.1’s English “accent” in syntax when generating Chinese outputs. Although Llama-3.1 is overall closer to the human distribution than other LLMs (as shown in Table 1), its syntactic structures still exhibit greater similarity

to human-written English than to native Chinese, suggesting that English syntactic patterns influence its Chinese generations. We can observe a similar effect in French, even if less noticeable since English and French belong to linguistic families (West Germanic and Romance, respectively) that are closer. The English accent of Llama-3.1 can also be seen in our case study of unnatural language patterns in text generated by Llama-3.1 in Chinese. Table 2 presents examples of POS tag  $n$ -grams that frequently appear in both English and Llama-3.1-generated Chinese, but much less frequently in native Chinese.

**Impact of pretraining data.** Among the evaluated models, the Qwen series consistently underperforms in naturalness across all languages, including Chinese, despite expectations that it would perform better due to its development in China and likely greater access to Chinese pretraining data. In contrast, Llama consistently produces more natural outputs. We hypothesize that Qwen’s reliance on synthetic data during pretraining (Yang et al., 2024), in contrast to Llama’s explicit avoidance of synthetic data (Dubey et al., 2024), may explain this discrepancy. Notably, Mistral-v0.3 achieves the highest naturalness in English, possibly due to being the least multilingual among the models evaluated. Additionally, Mistral models show strong naturalness in French, likely benefiting from the use of more authentic French data, given that the team behind it is based in France. Unfortunately, we lack detailed information about the exact composition of the training data for any of these models. We hope to see future releases of multilingual LLMs following the principles of OLMo (Groeneveld et al., 2024), where not only the model weights but also the training data and pipeline are made open source.

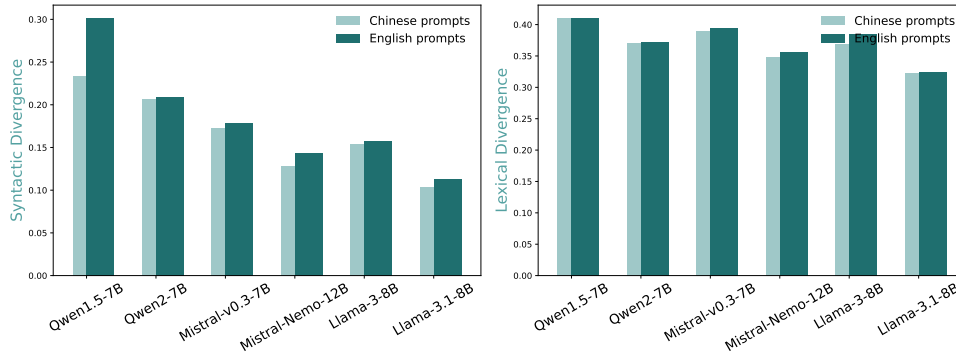


Figure 2: Comparison of divergence metrics for model generations in Chinese when the prompting language is either Chinese or English.

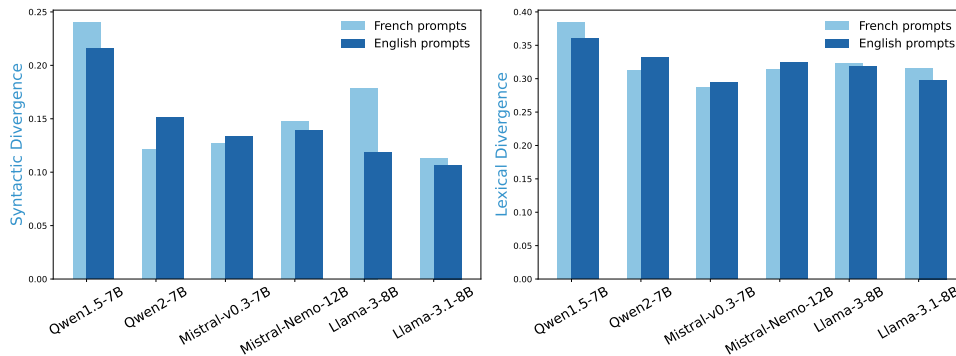


Figure 3: Comparison of divergence metrics for model generations in French when the prompting language is either French or English.

**Impact of prompting language.** For the previous benchmark results, we prompt the models in the same language in which we expect them to generate. To evaluate whether the input language affects the naturalness of the generations, we conduct an ablation study in a cross-lingual setting. Specifically, we manually translate the Chinese and French prompts into English and feed these English prompts to the LLM, while still instructing it to generate responses in Chinese or French. In other words, only the instructions are modified, the target language of the output remains unchanged. The results for the Chinese and French generations are presented in Figures 2 and 3. For Chinese outputs, prompting in Chinese consistently leads to better naturalness. However, for French outputs, the results vary across models. We hypothesize that this is influenced by the proportion of instruction-response pairs the models encountered in different languages during instruction tuning. Additionally, the closer typological relationship between French and English, compared to Chinese, may also contribute.

**Impact of decoding temperature.** Previous benchmark results are generated using temperature sampling with  $t = 0.6$ , which is found to provide the optimal balance between creativity and consistency for multilingual generation tasks (Dubey et al., 2024). As an ablation study, we vary the decoding temperature between 0.3 and 0.9 to examine its impact on the naturalness of model outputs. Our results in Figure 4 demonstrate that for syntactic naturalness, models that are already natural become even better as the decoding temperature increases, whereas models with unnatural outputs show a further increase in unnaturalness. For vocabulary, all models show improved naturalness as the decoding temperature rises. It is also worth noting that the ranking of the models generally remains unchanged across different decoding settings.

## 5 Improving Naturalness through Preference Tuning

We now propose an approach to improve the naturalness of language model outputs. Among the various stages of LLM development, preference-based learning is the most effective for refining

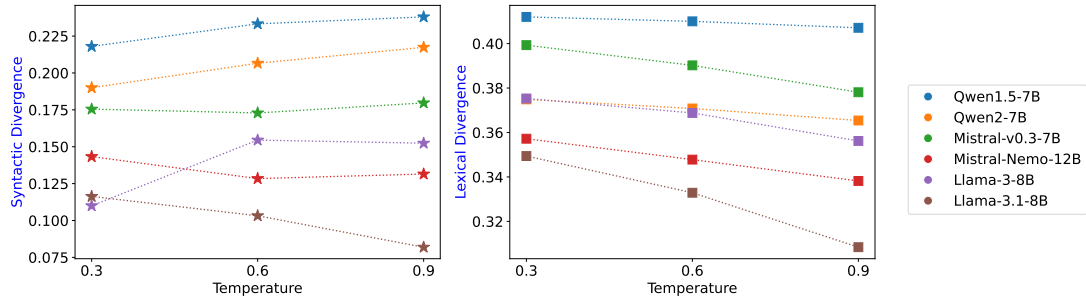


Figure 4: Variation of divergence metrics for model generations in Chinese when the decoding temperature changes.

		Llama-3.1-8B			Qwen2-7B		
		Unaligned	Llama data	Qwen data	Unaligned	Llama data	Qwen data
<b>Essay Generation</b>	Lexical Divergence ↓	34.89	35.06	<b>34.56</b>	41.45	<b>40.58</b>	40.60
	Syntactic Divergence ↓	12.63	12.59	<b>11.75</b>	22.66	<b>21.46</b>	22.41
	CMMLU Accuracy ↑	55.38	<b>55.62</b>	55.61	80.08	80.29	<b>80.36</b>
<b>Question Answering</b>	Lexical Divergence ↓	28.05	<b>28.01</b>	28.04	28.09	27.58	<b>26.92</b>
	Syntactic Divergence ↓	12.19	<b>9.44</b>	9.79	12.25	<b>10.78</b>	11.19
	CMMLU Accuracy ↑	55.38	<b>55.47</b>	<b>55.47</b>	80.08	80.47	<b>80.49</b>

Table 3: Naturalness alignment results for essay generation and question answering tasks in Chinese. Lower divergence values indicate better naturalness, while higher CMMLU accuracy reflects stronger performance.

stylistic features (Iverson et al., 2024). Based on this, we focus on aligning models for better naturalness during the preference tuning stage. Among the available preference optimization methods, we opt for DPO due to its simplicity and efficiency.

### 5.1 Preference Dataset Construction

We construct the preference datasets starting with SFT (Supervised Fine-Tuning) datasets originally written in the target language. However, we find that non-translated, open-source SFT datasets for non-English languages are almost non-existent. For French, we are unable to locate any such datasets, and for Chinese – despite being the second highest-resource language after English – only a limited number is available. Consequently, we conduct our DPO experiments on Chinese, using two datasets that focus on essay generation and open-domain QA tasks. These tasks are well-suited for stylistic alignment due to their open-ended, creative nature. For essay generation, we use the “composition” split from the Firefly dataset (Yang, 2023), and for open-domain QA, we use the OpenLabel-Chinese Conversations Dataset<sup>5</sup> (BAAI, 2023).

The original instructions from the SFT datasets are preserved, with the initial response serving as the preferred response. To generate the unnatural, i.e., rejected, response we apply synthetic manipu-

lations through paraphrasing and back-translation via English. The prompts we use to generate the rejected responses are listed in Appendix E. We ensure semantic consistency by using BLEU (Papineni et al., 2002) to filter out pairs of chosen and rejected responses with insufficient similarity. Moreover, to ensure that the linguistic style of the rejected responses differs significantly from the chosen ones, we also filter out pairs with overly high BLEU scores. Based on an empirical analysis of BLEU score distributions, we define the final threshold as  $0.15 < \text{BLEU}(\text{Chosen}, \text{Rejected}) < 0.9$ . We also filter out responses shorter than 10 words (often in the style of multiple-choice questions), as they do not convey enough lexical or syntactic style. The statistics of our two resulting preference datasets are presented in Table 8.

### 5.2 Experimental Setup

We believe that it is beneficial to perform naturalness alignment on models that have already undergone SFT and previous rounds of preference tuning, since preference tuning is typically done iteratively (Dubey et al., 2024). We conduct experiments with the instruction-tuned versions of Llama-3.1 and Qwen2, which demonstrated the highest and lowest performance, respectively, in naturalness for Chinese in Table 1. To increase GPU memory usage efficiency and to avoid catas-

<sup>5</sup>Released under the Apache 2.0 license.



Theme of Prompt	Output
夹乒乓大赛	Unaligned 今天, 老师组织了一场特殊的乒乓球比赛——夹乒乓大赛。<.....> 最后, 比赛结束了, 胜利者被宣布出来。 <i>Today, the teacher organized a special ping-pong competition—the Ping-Pong Clamping Contest. &lt;.....&gt; In the end, the competition concluded, and the winner was announced.</i>
The Ping-Pong Clamping Contest	Aligned 今天, 老师突然宣布要举办一场夹乒乓大赛。<.....> 最后, 夹得最快的人赢得了夹乒乓大赛的冠军。 <i>Today, the teacher suddenly announced that there would be a Ping-Pong Clamping Contest. &lt;.....&gt; In the end, the person who clamped the fastest won the championship of the Ping-Pong Clamping Contest.</i>

Table 4: Example of improved output from Llama-3.1 after applying naturalness alignment for Chinese essay generation. English translations are presented in italics. Unnatural expressions are marked in orange, while natural alternatives are highlighted in blue. Although the expression “the winner was announced” is common in English, its literal translation in Chinese is rarely used, despite being grammatically correct. It uses a passive construction which is far less common in Chinese and sounds overly formal.

trophic forgetting, we apply Low-Rank Adaptation (LoRA) (Hu et al., 2022) in conjunction with DPO. Detailed hyperparameters for both LoRA and DPO are provided in Appendix F.

We explore two approaches: (1) self-alignment, where the preference-tuning dataset is generated by the same model being aligned (e.g., using Llama-3.1 to create training data for Llama-3.1), and (2) cross-model alignment, where one model generates the dataset used to align a different model (e.g., using Llama-3.1 to produce training data for Qwen2).

### 5.3 Results

The results in Table 3 show consistent improvements in both lexical and syntactic naturalness across models and tasks after applying our alignment method. In Table 4, we provide an example on the essay generation task, comparing responses generated by Llama-3.1 before and after naturalness alignment using the same prompts.

Although we experiment with generating rejected responses using both paraphrasing and back-translation, the latter consistently yields slightly better performance than the former. We believe this is due to back-translation being more effective at introducing translationese artifacts into the rejected responses. Therefore, the final results in Table 3 are based on back-translation. However, there is no definitive conclusion on whether rejected responses should be generated using self-alignment or cross-model alignment, and future work could explore generating response pairs with a combination of different models.

In addition, we evaluate our aligned models on the Chinese Massive Multitask Language Understanding (CMMLU) benchmark (Li et al., 2023) to ensure that their general capabilities are not compromised after naturalness alignment. Our results show that our alignment method not only improves linguistic naturalness but also slightly enhances

overall language understanding performance. Interestingly, we observe that naturalness does not always positively correlate with language understanding performance. For example, Qwen2, despite achieving much higher scores than Llama-3.1 on the CMMLU benchmark, demonstrates lower naturalness, potentially due to the heavy use of synthetic data. This highlights the need for future research to consider naturalness as a complementary metric alongside conventional benchmark scores.

## 6 Conclusion

In this paper, we address the naturalness challenges of LLMs, particularly in multilingual contexts, by conducting experiments in English, French and Chinese. We introduce two corpus-level metrics to quantify the naturalness of model output distributions: one focused on vocabulary (lexical naturalness) and the other on grammatical structure (syntactic naturalness). These metrics are interpretable and free from biases introduced by external embeddings. Using these metrics, we benchmark state-of-the-art multilingual LLMs and analyze how factors such as training data, prompting language, and decoding strategy influence the generated language. Our analysis provides valuable insights into the current multilingual LLM landscape, complementing traditional performance benchmarks for task-solving capabilities. Finally, we propose an alignment method using DPO and a synthetically manipulated preference dataset to enhance the naturalness of model outputs. Experiments show that our aligned models consistently improve in both lexical and syntactic naturalness.

### Limitations

Our experiments do not include any low-resource languages because our approach relies on the availability of ground-truth distributions from native,

human-written data, which is often scarce or unavailable for many languages. While our naturalness evaluation metrics are designed to be language-agnostic, they depend on reliable word tokenizers and dependency parsers for the languages studied. This is discussed in greater detail in Appendix G. Unfortunately, such tools are still lacking for most low-resource languages. However, we argue that naturalness evaluation should not be the priority for these languages at this stage, as it is only meaningful once models achieve a certain baseline level of overall performance.

Our cross-lingual benchmark was limited to the Wikipedia domain, as this domain provides topically aligned, natively written content across languages. Wikipedia was the only source we could access with the necessary data. Although Wikipedia text serves as a reasonable proxy for general domain text, it may not guarantee that our findings are applicable across other domains. In the future, it would be valuable to develop more non-translated cross-lingual corpora for additional domains and extend the naturalness evaluation to those areas as well.

Our alignment approach has so far only been tested on data from essay generation and general domain question-answering tasks. These tasks were chosen because their creative and open-ended nature allows for the expression of stylistic features. However, applying this method to more knowledge-intensive and constrained tasks may introduce unintended knowledge editing, potentially leading to increased hallucination risks. Furthermore, our experiments were limited to Chinese due to the lack of natively written SFT datasets in other non-English languages. The Aya dataset (Singh et al., 2024), while a valuable resource for multilingual instruction tuning with native data, provides too few samples in each language. For example, after filtering for response length, we obtained only 958 samples in French from Aya, which is insufficient for our alignment and evaluation approach.

Collecting human annotations for naturalness evaluation is challenging. We initially attempted to gather human annotations for a meta-evaluation of our metrics, but annotators reported difficulty in distinguishing linguistic naturalness from individual generations. Our naturalness evaluation operates on a corpus level. As discussed in Section 3, while a single text with specific vocabulary choices or grammatical structures may appear natural, repeated occurrences of these features across many

generations would raise concerns. Human evaluators cannot easily process a large corpus and identify these patterns. Therefore, we believe that our automatic metrics address this gap where human evaluations fall short.

Finally, our evaluation and alignment methods focus solely on the naturalness of linguistic form, without considering social biases. However, we believe linguistic biases are significantly underexplored compared to social biases, and we aim to bridge this gap by taking an initial step in this direction.

## Acknowledgments

We would like to thank all the people at Apple who supported this work and provided helpful feedback. The majority of this work was carried out during Yanzhu Guo’s internship at Apple. Simone Conia gratefully acknowledges the support of the PNRR MUR project PE0000013-FAIR, which fully funds his fellowship since October 2023.

## References

- Marah Abidin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.
- BAAI. 2023. Openlabel-chinese conversations dataset. <https://data.baai.ac.cn/details/OL-CC>.
- Yuri Bizzoni, Tom S Juzek, Cristina España-Bonet, Koel Dutta Chowdhury, Josef van Genabith, and Elke Teich. 2020. [How human is machine translationese? comparing human and machine translations of text and speech](#). In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 280–290, Online. Association for Computational Linguistics.
- Ilias Chalkidis, Tommaso Pasini, Sheng Zhang, Letizia Tomada, Sebastian Schwemer, and Anders Søgaard. 2022. [FairLex: A multilingual benchmark for evaluating fairness in legal text processing](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4389–4406, Dublin, Ireland. Association for Computational Linguistics.
- Pinzhen Chen, Zhicheng Guo, Barry Haddow, and Kenneth Heafield. 2024. [Iterative translation refinement with large language models](#). In *Proceedings of the 25th Annual Conference of the European Association for Machine Translation (Volume 1)*, pages 181–190, Sheffield, UK. European Association for Machine Translation (EAMT).

- Susanne DeVore and Kristopher Kyle. 2023. Assessing syntactic and lexicogrammatical use in second language mandarin writing samples. *Journal of Second Language Writing*, 60:101014.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Shangbin Feng, Vidhisha Balachandran, Yuyang Bai, and Yulia Tsvetkov. 2023. [FactKB: Generalizable factuality evaluation using language models enhanced with factual knowledge](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 933–952, Singapore. Association for Computational Linguistics.
- Markus Freitag, David Vilar, David Grangier, Colin Cherry, and George Foster. 2022. [A natural diet: Towards improving naturalness of machine translation output](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3340–3353, Dublin, Ireland. Association for Computational Linguistics.
- Mario Giulianelli, Joris Baan, Wilker Aziz, Raquel Fernández, and Barbara Plank. 2023. [What comes next? evaluating uncertainty in neural text generators against human production variability](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14349–14371, Singapore. Association for Computational Linguistics.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. 2012. [A kernel two-sample test](#). *Journal of Machine Learning Research*, 13(25):723–773.
- Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, et al. 2024. Olmo: Accelerating the science of language models. *arXiv preprint arXiv:2402.00838*.
- Isabel Groves, Ye Tian, and Ioannis Douratsos. 2018. [Treat the system like a human student: Automatic naturalness evaluation of generated text without reference texts](#). In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 109–118, Tilburg University, The Netherlands. Association for Computational Linguistics.
- Biyang Guo, Xin Zhang, Ziyuan Wang, Minqi Jiang, Jinran Nie, Yuxuan Ding, Jianwei Yue, and Yupeng Wu. 2023. How close is chatgpt to human experts? comparison corpus, evaluation, and detection. *arXiv preprint arxiv:2301.07597*.
- Yanzhu Guo, Guokan Shang, Michalis Vazirgiannis, and Chloé Clavel. 2024. [The curious decline of linguistic diversity: Training language models on synthetic text](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3589–3604, Mexico City, Mexico. Association for Computational Linguistics.
- Melissa Heikkilä. 2023. [How ai-generated text is poisoning the internet](#). *MIT Technology Review*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Hamish Ivison, Yizhong Wang, Jiacheng Liu, Zeqiu Wu, Valentina Pyatkin, Nathan Lambert, Noah A Smith, Yejin Choi, and Hannaneh Hajishirzi. 2024. [Unpacking dpo and ppo: Disentangling best practices for learning from preference feedback](#). *arXiv preprint arXiv:2406.09279*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Haonan Li, Yixuan Zhang, Fajri Koto, Yifei Yang, Hai Zhao, Yeyun Gong, Nan Duan, and Timothy Baldwin. 2023. [Cmmlu: Measuring massive multitask language understanding in chinese](#). *arXiv preprint arXiv:2306.09212*.
- Weixin Liang, Zachary Izzo, Yaohui Zhang, Haley Lepp, Hancheng Cao, Xuandong Zhao, Lingjiao Chen, Haotian Ye, Sheng Liu, Zhi Huang, Daniel McFarland, and James Y. Zou. 2024a. [Monitoring AI-modified content at scale: A case study on the impact of chatGPT on AI conference peer reviews](#). In *Forty-first International Conference on Machine Learning*.
- Weixin Liang, Yaohui Zhang, Zhengxuan Wu, Haley Lepp, Wenlong Ji, Xuandong Zhao, Hancheng Cao, Sheng Liu, Siyu He, Zhi Huang, Diyi Yang, Christopher Potts, Christopher D Manning, and James Y. Zou. 2024b. [Mapping the increasing use of LLMs in scientific papers](#). In *First Conference on Language Modeling*.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Ye Liu, Wolfgang Maier, Wolfgang Minker, and Stefan Ultes. 2021. [Naturalness evaluation of natural language generation in task-oriented dialogues using BERT](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 839–845, Held Online. INCOMA Ltd.

- Jiaming Luo, Colin Cherry, and George Foster. 2024. [To diverge or not to diverge: A morphosyntactic perspective on machine translation vs human translation](#). *Transactions of the Association for Computational Linguistics*, 12:355–371.
- Clara Meister and Ryan Cotterell. 2021. [Language model evaluation beyond perplexity](#). 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 5328–5339, Online. Association for Computational Linguistics.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D. Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. [Universal Dependencies v2: An evergrowing multilingual treebank collection](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4034–4043, Marseille, France. European Language Resources Association.
- Jekaterina Novikova, Oliver Lemon, and Verena Rieser. 2016. [Crowd-sourcing NLG data: Pictures elicit better data](#). In *Proceedings of the 9th International Natural Language Generation conference*, pages 265–273, Edinburgh, UK. Association for Computational Linguistics.
- Isabel Papadimitriou, Kezia Lopez, and Dan Jurafsky. 2023. [Multilingual BERT has an accent: Evaluating English influences on fluency in multilingual models](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1194–1200, Dubrovnik, Croatia. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Andrew J Peterson. 2024. [Ai and the problem of knowledge collapse](#). *arXiv preprint arXiv:2404.03502*.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. 2021. [Mauve: Measuring the gap between neural text and human text using divergence frontiers](#). *Advances in Neural Information Processing Systems*, 34:4816–4828.
- Tiago Pimentel, Clara Isabel Meister, and Ryan Cotterell. 2023. [On the usefulness of embeddings, clusters and strings for text generation evaluation](#). In *The Eleventh International Conference on Learning Representations*.
- Barbara Plank. 2022. [The “problem” of human label variation: On ground truth in data, modeling and evaluation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10671–10682, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online. Association for Computational Linguistics.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Nino Shervashidze, Pascal Schweitzer, Erik Jan Van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. 2011. [Weisfeiler-lehman graph kernels](#). *Journal of Machine Learning Research*, 12(9).
- Iliia Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal. 2024. [Ai models collapse when trained on recursively generated data](#). *Nature*, 631(8022):755–759.
- Giannis Siglidis, Giannis Nikolentzos, Stratis Limnios, Christos Giatsidis, Konstantinos Skianis, and Michalis Vazirgiannis. 2020. [Grakel: A graph kernel library in python](#). *Journal of Machine Learning Research*, 21(54):1–5.
- Shivalika Singh, Freddie Vargus, Daniel D’souza, Börje Karlsson, Abinaya Mahendiran, Wei-Yin Ko, Herumb Shandilya, Jay Patel, Deividas Mataciunas, Laura O’Mahony, Mike Zhang, Ramith Het-tiarachchi, Joseph Wilson, Marina Machado, Luisa Moura, Dominik Krzemiński, Hakimeh Fadaei, Irem Ergun, Ifeoma Okoh, Aisha Alaagib, Oshan Mudannayake, Zaid Alyafeai, Vu Chien, Sebastian Ruder, Surya Guthikonda, Emad Alghamdi, Sebastian Gehrmann, Niklas Muennighoff, Max Bartolo, Julia Kreutzer, Ahmet Üstün, Marzieh Fadaee, and Sara Hooker. 2024. [Aya dataset: An open-access collection for multilingual instruction tuning](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11521–11567, Bangkok, Thailand. Association for Computational Linguistics.
- Guy Tevet and Jonathan Berant. 2021. [Evaluating the evaluation of diversity in natural language generation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 326–346, Online. Association for Computational Linguistics.
- Eva Vanmassenhove, Dimitar Shterionov, and Matthew Gwilliam. 2021. [Machine translationese: Effects of algorithmic bias on linguistic complexity in machine translation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2203–2213, Online. Association for Computational Linguistics.

- Vered Volansky, Noam Ordan, and Shuly Wintner. 2015. On the features of translationese. *Digital Scholarship in the Humanities*, 30(1):98–118.
- Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. 2024. [Do llamas work in English? on the latent language of multilingual transformers](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15366–15394, Bangkok, Thailand. Association for Computational Linguistics.
- Shuly Wintner. 2016. [Translationese: Between human and machine translation](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Tutorial Abstracts*, pages 18–19, Osaka, Japan. The COLING 2016 Organizing Committee.
- 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.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- Jianxin Yang. 2023. [Firefly\(流萤\): 中文对话式大语言模型](#). <https://github.com/yangjianxin1/Firefly>.
- Tianyi Zhang\*, Varsha Kishore\*, Felix Wu\*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Zhexin Zhang, Leqi Lei, Lindong Wu, Rui Sun, Yongkang Huang, Chong Long, Xiao Liu, Xuanyu Lei, Jie Tang, and Minlie Huang. 2024. [SafetyBench: Evaluating the safety of large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15537–15553, Bangkok, Thailand. Association for Computational Linguistics.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhonghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging LLM-as-a-judge with MT-bench and chatbot arena](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Jürgen Ziegler and Tim Donkers. 2024. [From explanations to human-ai co-evolution: charting trajectories towards future user-centric ai](#). *i-com*, 23(2):263–272.

## A Preprocessing of the Wikipedia Dataset

We truncate any summaries exceeding 1024 tokens, as determined by Llama-3.1’s tokenizer. Since the Chinese version of Wikipedia often contains content in traditional Chinese, we use the zhconv-rs<sup>6</sup> library to convert all text to simplified Chinese for consistency during pre-processing.

## B Wikipedia Description Generation

We now present the employed prompts and settings for Wikipedia Description Generation.

### B.1 Prompts

The prompts used in our experiments are listed in Table 5. These prompts are then formatted using the default chat template<sup>7</sup> provided with each model in the Transformers library. We intentionally keep the prompts simple and do not include in-context learning to ensure the model’s behavior remains generalizable. After manually reviewing 50 generations from each model, we confirm that all models can follow the instructions and produce outputs aligned with our expectations. While the generated texts may not always be completely factual, they are consistently fluent and relevant to the given prompt, making them suitable for analyzing the general linguistic patterns of language models. Additionally, we use the fasttext-langdetect library<sup>8</sup> to identify the language of the generations, filtering out those that are not in the target language. Across all models and languages, more than 99% of the generations are correctly classified in the target language.

### B.2 Generation Settings

We use bf16 precision for generation, with a max\_new\_tokens limit of 1024, matching the truncation length of the human-written summaries. For most experiments, except those analyzing the effect of decoding temperature, we use temperature sampling with a temperature of 0.6 and a repetition\_penalty of 1.02.

## C Implementation of Naturalness Metrics

For lexical naturalness, we process vocabulary at the word level instead of subword token level to

<sup>6</sup><https://github.com/Gowee/zhconv-rs>

<sup>7</sup>[https://huggingface.co/docs/transformers/main/en/chat\\_templating](https://huggingface.co/docs/transformers/main/en/chat_templating)

<sup>8</sup><https://pypi.org/project/fasttext-langdetect/>

Target Language	Prompt Language	Prompt
English	English	Please write a summary description of the following Wikipedia entry in English: <wiki entry in English>
Chinese	Chinese	请用简体中文写出以下维基百科词条的摘要描述: <wiki entry in Chinese>
	English	Please write a summary description of the following Wikipedia page in Simplified Chinese: <wiki entry in Chinese>
French	French	Veillez rédiger une description résumée de l'entrée Wikipedia suivante en français: <wiki entry in French>
	English	Please write a summary description of the following Wikipedia entry in French: <wiki entry in French>

Table 5: Prompts for Wikipedia description generation. For generations in Chinese and French, we experiment with both prompts in the target language and in English, to study the impact of prompting language on naturalness.

Transformation	Prompt
	You are a professional translator. You always show the translated version, without any additional explanations or format changes.
Back Translation	zh->en Translate from English into Simplified Chinese: <text in Chinese>
	You are a professional translator. You always show the translated version, without any additional explanations or format changes.
en->zh	Translate from Simplified Chinese into English: <text in English>
	You are a professional editor who revises word choices and restructures sentences while preserving the original meaning. You always show the edited version, without any additional explanations or format changes.
Paraphrasing	Edit the following text: <text in Chinese>

Table 6: Prompts used to generate rejected responses in the preference tuning dataset.

represent more meaningful lexical units. We use the Jieba<sup>9</sup> tokenizer for Chinese and the NLTK<sup>10</sup> tokenizer for French and English. We remove punctuation and digits, but do not discard stop words as they are an important linguistic feature (Meister and Cotterell, 2021).

For syntactic divergence, we parse sentences into dependency trees using the Stanza toolkit<sup>11</sup> (Qi et al., 2020), which generates dependency trees for each sentence according to the UD grammar (Nivre et al., 2020). We use the implementation of the WL kernel in the GraKeL library (Siglidis et al., 2020) and normalize all kernel values between 0 and 1.

The divergence values are calculated between large distributions (60K words for lexical divergence and 3K sentences for syntactic divergence), so the measures are relatively stable. We try bootstrapping with 10 different randomizations for each measure and find the variation interval to be within 5% for both.

<sup>9</sup><https://github.com/fxsjy/jieba>

<sup>10</sup><https://www.nltk.org>

<sup>11</sup>Released under the ODC-By v1.0 license.

## D Comparison with N-gram Overlap Metrics

We compute BLEU and ROUGE scores for the models evaluated in the naturalness benchmark (Table 1) and report them in Table 7. These n-gram overlap metrics show markedly different behavior from our naturalness metrics. For example, while our metrics consistently identify the Qwen models as the least natural across all three languages, this pattern is not reflected in their BLEU or ROUGE scores. Similarly, our metrics capture a steady improvement in naturalness in newer LLM versions compared to earlier ones, a trend that BLEU and ROUGE fail to detect.

## E Preference Dataset Generation

We provide the prompts used for generating our preference dataset, along with statistical insights into the resulting dataset.

### E.1 Prompts

The prompts used for generation of the preference dataset are presented in Table 6. Here, we only use English prompts, as our previous experiments show

		Qwen1.5-7B	Qwen2-7B	Mistral-v0.3-7B	Mistral-Nemo-12B	Llama-3-8B	Llama-3.1-8B
<b>English</b>	ROUGE-L	14.98	17.81	15.36	18.49	17.52	18.35
	BLEU	1.01	2.68	1.21	2.94	2.91	4.93
<b>Chinese</b>	ROUGE-L	11.56	13.65	13.17	10.81	13.14	12.16
	BLEU	0.56	1.05	0.93	0.48	0.96	0.99
<b>French</b>	ROUGE-L	13.28	16.88	16.56	16.14	16.78	15.88
	BLEU	0.67	2.40	2.69	1.91	3.58	3.25

Table 7: BLEU and ROUGE scores for the models evaluated in the naturalness benchmark.

	#Train	#Test	Len <sub>prompt</sub>	Len <sub>chosen</sub>	Len <sub>rejected</sub>	BLEU( <i>chosen, rejected</i> )
Essay	10,050	4,950	29	564	523	0.29
QA	4,281	2,109	52	185	176	0.46

Table 8: Statistics of our preference tuning datasets constructed with back translation. All tokenization for length calculation was performed using the Llama-3.1 tokenizer.

that prompting in English degrades the naturalness of generated Chinese, helping us produce the desired unnatural responses. The generation setting is the same as for Wikipedia Description Generation.

## E.2 Dataset Statistics

Statistics of the generated preference dataset is shown in Table 8. Rejected responses are shorter than the chosen ones on average, which may be due to models’ tendencies to prioritize words that occur more frequently in the tokenizer’s training set, resulting in fewer subword splits.

## F Hyperparameters for Preference Tuning

We conduct experiments using 8 Nvidia A100 GPUs, each with 40GB of memory. All models are trained in bf16 precision for 1 epoch. DPO training was performed with data parallelism, taking approximately 2 to 3 hours per model per dataset. We utilize the DPO implementation from the trl library<sup>12</sup> and the LORA implementation from the PEFT library<sup>13</sup>. The best-performing hyperparameters are listed in Table 9.

## G Adaptation of Naturalness Metrics for Other Languages

In our current approach, the lexical naturalness metric operates at the word level, as words are meaningful lexical units in the three languages (English, French and Chinese) analyzed in this study. However, for polysynthetic and agglutinative languages,

which feature richer morphological structures and sparser word distributions, morphemes may be more suitable as the basic lexical units. However, adapting our metrics for these languages would require morphological parsing, which might be challenging in some cases, especially for polysynthetic languages. For the syntactic naturalness metric, we use POS units as defined by the UD grammar. In the languages we analyzed, these units typically align with individual words. However, in English and French, there exist multi-word tokens where a single orthographic token corresponds to multiple syntactic words, a detail which we will clarify in the final version of our paper. Extending our syntactic metrics to polysynthetic and agglutinative languages would similarly require using the specific annotation frameworks provided by the UD schema for each language. For instance, in Turkish, the UD schema also divides orthographic tokens into syntactic words. We thank the reviewer for raising this interesting question. We believe that future efforts to extend our framework to additional languages should address not only the resource disparity between high-resource and low-resource languages but also integrate linguistically informed decisions to account for the diversity of typological features across languages.

<sup>12</sup><https://huggingface.co/docs/trl/index>

<sup>13</sup><https://huggingface.co/docs/peft/index>

	<b>General</b>				<b>LORA</b>			<b>DPO</b>	
<b>Parameter</b>	learning_rate	max_grad_norm	warmup_ratio	per_device_batch_size	lora_alpha	lora_dropout	r	target_modules	beta
<b>Value</b>	5e-6	0.3	0.1	6	128	0.05	256	all-linear	0.5

Table 9: Hyperparameters used for preference tuning. Initial values are based on recommendations from the DPO and LoRA papers, with minimal additional tuning applied.