

No Language Data Left Behind: A Cross-Cultural Study of CJK Language Datasets in the Hugging Face Ecosystem

Dasol Choi^{1,2,3*} Woomyoung Park⁴ Youngsook Song^{5†}

¹Yonsei University ³AIM INTELLIGENCE ²MODULABS ⁴SAIONIC AI ⁵Lablup Inc.
dasolchoi@yonsei.ac.kr max@sionic.ai yssong@lablup.com

Abstract

Recent advances in Natural Language Processing (NLP) have underscored the crucial role of high-quality datasets in building large language models (LLMs). However, while extensive resources and analyses exist for English, the landscape for East Asian languages, particularly Chinese, Japanese, and Korean (CJK), remains fragmented and underexplored, despite these languages serving over 1.6 billion speakers. To address this gap, we investigate the HuggingFace ecosystem from a cross-linguistic perspective, focusing on how cultural norms, research environments, and institutional practices shape dataset availability and quality. Drawing on more than 3,300 datasets, we employ quantitative and qualitative methods to examine how these factors drive distinct creation and curation patterns across Chinese, Japanese, and Korean NLP communities. Our findings highlight the large-scale and often institution-driven nature of Chinese datasets, grassroots community-led development in Korean NLP, and an entertainment and subculture-focused emphasis on Japanese collections. By uncovering these patterns, we reveal practical strategies for enhancing dataset documentation, licensing clarity, and cross-lingual resource sharing, guiding more effective and culturally attuned LLM development in East Asia. We conclude by discussing best practices for future dataset curation and collaboration, aiming to strengthen resource development across all three languages.

1 Introduction

With the emergence of Large Language Models (LLMs) transforming the field of Natural Language Processing (NLP) (Kenton and Toutanova, 2019; Brown et al., 2020; Achiam et al., 2023), the importance of high-quality datasets in model development has become increasingly critical. For datasets to be valuable in this context, they must meet both

quantitative requirements (sufficient size and coverage) and qualitative standards (reliability and representativeness). While English-language resources have been extensively studied, the landscape of datasets for East Asian languages—particularly Chinese, Japanese, and Korean (CJK)—remains comparatively underexplored (Joshi et al., 2020; Bender, 2019). This gap is especially noteworthy given that these languages collectively serve over 1.6 billion speakers and originate from major hubs of technological innovation.

In recent years, platforms such as HuggingFace have emerged as central repositories for distributing and accessing NLP datasets, making these resources widely accessible while introducing new challenges in dataset discovery, quality assessment, and cross-lingual collaboration (Hugging Face, 2023; Lhoest et al., 2021). These challenges are particularly pronounced for CJK languages due to their unique linguistic features, distinct cultural contexts, and varying approaches to data sharing and documentation.

While CJK languages play an increasingly important role in global NLP research, several critical issues need to be addressed. First, there is a limited understanding of how dataset creation patterns differ across these language communities—and how those differences reflect their respective NLP ecosystems. Second, although cultural and institutional factors evidently influence dataset generation, their specific impacts on dataset characteristics and quality have yet to be systematically investigated. Third, the potential for cross-lingual synergies among CJK languages remains largely untapped, despite their many shared cultural and linguistic foundations. For instance, while benchmarks like MMLU (Hendrycks et al., 2020) are being generated in multiple languages, insufficient comparative analysis exists regarding how these variants differ from traditional parallel corpora or how download patterns across regions reflect dis-

*Work done during internship at SAIONIC AI.

†Corresponding author.

tinct cultural preferences. To address these challenges, this study investigates the HuggingFace ecosystem for Chinese, Japanese, and Korean NLP resources, focusing on dataset development and usage patterns. The specific objectives of this work are to:

- Examine the current landscape of CJK datasets on HuggingFace, including key metadata such as domain, dataset size, documentation practices, and usage statistics.
- Analyze the cultural characteristics and development patterns underlying dataset creation in each language community, highlighting commonalities and differences in curation and documentation.
- Identify potential synergies and cross-lingual opportunities among the three languages, and propose strategies for more effective, collaborative dataset development.

2 Related Work

Documentation efforts for CJK language resources have expanded over time. For Chinese, [Tao et al. \(2009\)](#) introduce systematic approaches for constructing and evaluating linguistic data resources, while [Li et al. \(2023a\)](#) compile a comprehensive repository emphasizing accessibility and classification frameworks. In the Korean context, [Cho et al. \(2020\)](#) survey existing NLP resources, and [Cho et al. \(2023\)](#) examine how local research cultures influence resource development. However, these studies largely predate large generative models and rely on GitHub as the primary hub, while systematic analyses of CJK datasets within newer ecosystems like HuggingFace remain rare.

Major multilingual projects including BigScience ([Le Scao et al., 2023](#)), CC100 ([Wenzek et al., 2019](#)), and LAION ([Schuhmann et al., 2022](#)) expand non-English data availability, yet don't investigate cultural factors affecting dataset usage within communities. The movement toward improved dataset documentation ([Geburu et al., 2021](#)) highlights ethical considerations, though adoption varies across languages and platforms.

HuggingFace has emerged as a central repository for NLP resources, but studies flag challenges including inconsistent documentation ([Yang et al., 2024](#)), limited transparency ([Pepe et al., 2024](#)), and ambiguous licensing. These issues underscore the

need to examine dataset curation in cultural contexts ([Lhoest et al., 2021](#)). Recent work by [Dargis et al. \(2024\)](#) offers insights for building evaluation frameworks for languages with particular traits, yet no study comprehensively analyzes how CJK dataset ecosystems are shaped by local research cultures, licensing preferences, and community-driven development.

Building on these perspectives, our research provides the first large-scale, comparative analysis of CJK datasets within the HuggingFace ecosystem, investigating how cultural contexts and documentation standards influence dataset usage patterns and resource quality.

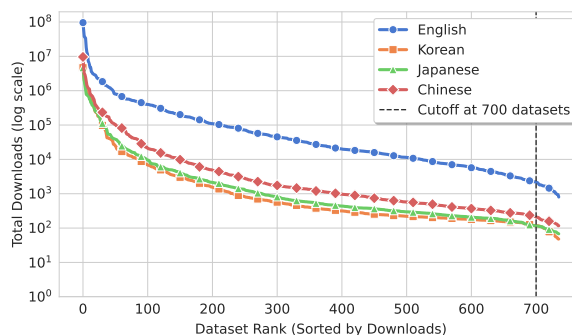


Figure 1: Datasets of each language sorted by number of downloads in descending order. Based on the decreasing pattern of downloads, we set the cutoff point at 700.

3 Method

3.1 Data Collection

We systematically collected dataset information from the HuggingFace platform using their Datasets API. Our data collection strategy focused on identifying actively used datasets for each target language (Chinese, Japanese, Korean, and English as a reference).

To determine a cutoff point for dataset inclusion, we analyzed the download frequency distribution for each language (Figure 1). The download counts follow a power-law distribution, with all languages showing consistent patterns. The distributions converge around the 700th dataset, where download counts fall below 100. Beyond this point, we observe minimal engagement and declining documentation quality. This natural boundary led us to set our cutoff at 700 datasets per language, ensuring both coverage and quality.

Datasets were retrieved from <https://huggingface.co/datasets> using language filters and download frequency sorting. All statistics were recorded on January 28, 2025.

Category	Evaluation Metrics
Scale & Composition	Dataset Size: Distribution across size categories (small, medium, large, extra-large) Language Makeup: Distribution of monolingual, English-paired, and multilingual datasets Task Types: Distribution of major NLP tasks (text generation, QA, classification, etc.)
Development Patterns	Ownership Structure: Proportions of corporate, institutional, and individual contributions License Types: Distribution of permissive, copyleft, unknown, and other licenses Community Activity: Dataset creation trends and instruction tuning development
Documentation Quality	Academic Validation: Presence of associated arXiv papers and research citations Documentation Standards: Adherence to HuggingFace dataset card templates Documentation Depth: Comprehensiveness of dataset descriptions and README files
Cultural Characteristics	Domain Focus: Specialized fields (e.g., medical, entertainment, content moderation) Resource Development: Approaches to dataset creation and curation Community Priorities: Language-specific preferences and development patterns

Table 1: Analysis framework for CJK datasets, organized by category.

For each dataset, we extracted metadata across four categories: **Scale & Composition** (dataset size, language combinations, task types, temporal patterns), **Development Patterns** (ownership, licensing, community metrics), **Documentation Quality** (dataset cards, citations, README files, metadata completeness), and **Cultural Characteristics** (domain focus, development approaches, community patterns).

In addition, we collected the complete dataset cards to analyze documentation practices and cultural characteristics in depth. We will release the full metadata and dataset card contents as a public resource upon publication.

3.2 Analysis Framework

Our analysis framework combines quantitative and qualitative approaches to examine CJK language datasets. Table 1 presents our analysis metrics across four main categories. For quantitative analysis, we focus on measuring dataset sizes, language distributions, task type proportions, ownership ratios, and documentation completeness. Our qualitative analysis examines domain preferences, resource development approaches, and community characteristics. This mixed-method approach helps us understand how dataset development patterns reflect each language community’s unique characteristics, particularly in terms of instruction tuning trends, domain preferences, and resource development strategies.

4 Results and Analysis

4.1 Overview of CJK Datasets


The stacked bar chart (Figure 2b) provides a view of the top 700 most downloaded datasets for each language, categorizing them into monolin-

gual, English-paired (bilingual with English), and multilingual (three or more languages) resources. In this broader analysis, Chinese datasets show the highest proportion of English-paired resources (148 datasets), notably higher than Korean (57) or Japanese (67), suggesting a greater emphasis on cross-lingual applications. The multilingual category shows substantial representation across all three languages, with similar proportions (Korean: 272, Japanese: 299, Chinese: 291), indicating active participation in multilingual resource development. While multilingual resources reflect diverse aspects of CJK datasets, we focus our subsequent analyses on *monolingual datasets* to better understand language-specific characteristics.

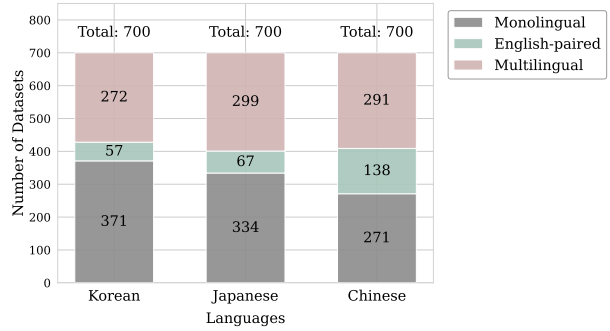
Dataset Size Distribution. Table 2 presents the size distribution of monolingual datasets across languages, categorized as Small (<10MB), Medium (10MB–100MB), Large (100MB–1GB), and Extra-Large (>1GB). Chinese datasets show a relatively balanced distribution across size categories, with a notable presence in large (10) and extra-large (12) categories. Japanese datasets demonstrate strong representation in small-scale resources (123) but notably lack extra-large datasets. Korean shows similar concentration in small datasets (137) and limited presence in medium and large categories, yet maintains a notable presence in the extra-large category (7). English datasets maintain the highest counts across all categories, providing a reference point for resource availability.

4.2 Comparative Characteristics of CJK Datasets

Task Distribution Figure 3 presents the distribution of task categories across languages through a



(a) Distribution of CJK Language Intersection



(b) Composition of Language Types in Datasets

Figure 2: Distribution and Composition Analysis of CJK Language Datasets. (a) Illustrates the intersections among CJK language datasets, showing unique and overlapping dataset counts. (b) Shows the composition of the top 700 downloaded datasets for each language, categorized into monolingual, English-paired, and multilingual resources.

Size	English	Chinese	Japanese	Korean
S (<10M)	258	144	123	137
M (10M–100M)	21	11	3	3
L (100M–1B)	14	10	4	3
XL (>1B)	22	12	-	7

Table 2: Dataset size distribution across languages.

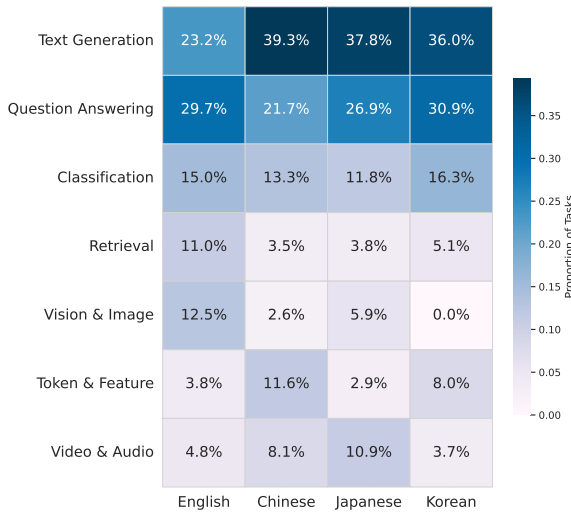


Figure 3: Task distribution across different languages. The heatmap illustrates the proportion of datasets belonging to the top 7 most frequent task categories across English, Chinese, Japanese, and Korean datasets.

heatmap visualization. *Text Generation* emerges as the dominant task across all languages, with particularly high proportions in Chinese (39.3%),

Task categories group related NLP tasks based on their functionality (e.g., Text Generation includes text-generation, language-modeling, fill-mask; Question Answering includes question-answering, multiple-choice, extractive-qa).

Japanese (37.8%), and Korean (36.0%) datasets compared to English (23.2%). *Question Answering* follows as the second most common task, with moderate variations: Korean (30.9%) and English (29.7%) show higher proportions than Chinese (21.7%) and Japanese (26.9%).

The analysis reveals distinctive task preferences across languages beyond these two categories. *Classification* tasks appear between 11.8% and 16.3% of datasets, with Korean (16.3%) having the highest ratio and Japanese (11.8%) the lowest. *Token & Feature* tasks are more prominent in Chinese (11.6%) and Korean (8.0%) than in English (3.8%) or Japanese (2.9%). *Video & Audio* tasks show varied representation, with Japanese leading at 10.9%, followed by Chinese (8.1%), English (4.8%), and Korean (3.7%). Lastly, *Vision & Image* tasks exhibit particularly striking differences: English leads at 12.5%, followed by Japanese (5.9%) and Chinese (2.6%), while Korean shows no representation (0.0%). *Retrieval* tasks also show notable variation, with English (11.0%) significantly ahead of other languages (3.5-5.1%).

These distinctive patterns reflect differing research priorities and technological needs across language communities. The Japanese emphasis on *Video & Audio* may correspond to its strong anime and entertainment industry, while Korean’s focus on *Classification* and Chinese’s on *Token & Feature* tasks suggest prioritization of fundamental NLP infrastructure development tailored to their respective linguistic complexities.

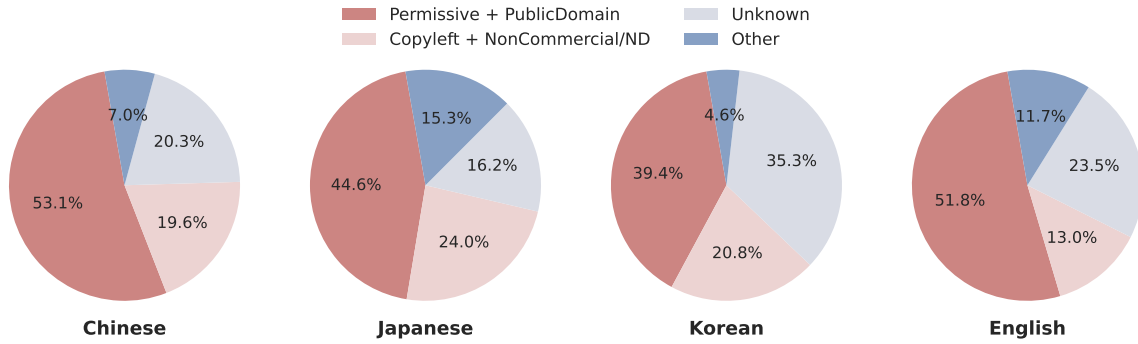


Figure 4: License distribution across CJK and English datasets, showing the proportion of *Permissive + PublicDomain*, *Copyleft + NonCommercial/ND*, *Unknown*, and *Other* licenses for each language community.




Figure 5: Instruction Datasets Over Time by Language (English and CJK), from late 2022 to 2024.

Evolution of Instruction Tuning Analysis of instruction-tuning datasets reveals distinct patterns across languages. Among the top 700 most downloaded datasets for each language, instruction-tuning datasets show notable presence: Korean (13.3%), Japanese (12.6%), Chinese (12.3%), and English (7.0%). This higher proportion in CJK languages compared to English suggests particularly active instruction-tuning development in these communities.

Analysis of temporal patterns (Figure 5) reveals distinct characteristics across languages. Chinese datasets show notable early activity in 2023, peaking around 7 releases. Korean datasets demonstrate dramatic fluctuations in 2024, reaching highest peaks of 12-14 releases mid-2024. Japanese datasets show moderate initial activity but increased activity during 2024, reaching peaks of 7-9 releases. English datasets maintain stable patterns throughout, typically with 1-4 releases monthly.

These patterns reflect different community approaches to instruction dataset development:

- **Chinese:** Early adoption with moderate peaks

We identify instruction datasets through ‘instruct’ keywords and common dataset names in metadata.

(around 7 releases) followed by decreased activity

- **Korean:** Shows the highest peaks (up to 14 releases) with considerable volatility
- **Japanese:** Late but substantial increase in development activity
- **English:** Consistent but moderate release patterns throughout

License Distribution Patterns Figure 4 shows how each language community approaches data licensing. Two major observations emerge. First, Chinese (53.1%) and English (51.8%) exhibit the highest proportions of *Permissive* or *Public Domain* licenses, indicating a shared culture of open access. Both also include moderate segments of *Unknown* (Chinese 20.3%, English 23.5%) and *Copyleft/NonCommercial* (19.6% and 13.0% respectively), suggesting a balance between openness and controlled usage.

Second, Japanese (44.6%) and Korean (39.4%) have lower shares of *Permissive/Public Domain* compared to Chinese and English but differ substantially in other categories. Japanese devotes 24.0% to *Copyleft/NonCommercial*—more than any other language—while also having a relatively high *Other* category (15.3%). In contrast, Korean stands out for its large *Unknown* portion (35.3%), underscoring possible gaps in documentation practices despite still having a notable *Copyleft/NonCommercial* share (20.8%). Taken together,

Licenses were classified into five categories: *Permissive* (e.g., Apache, MIT, CC-BY), *Public Domain* (e.g., CC0, PDDL), *Copyleft/ShareAlike* (e.g., GPL, CC-BY-SA), *Non-Commercial/ND* (CC-BY-NC, CC-BY-ND), and *Other*. For visualization, *Permissive* and *Public Domain* categories were combined, as were *Copyleft/ShareAlike* and *Non-Commercial/ND*.

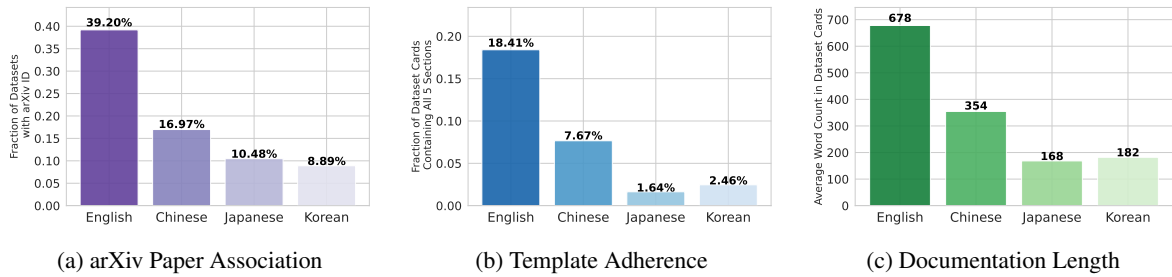


Figure 6: Comparison of Documentation Quality Across Languages. (a) Shows the percentage of datasets associated with academic publications, (b) presents the percentage of dataset cards containing all five structured sections, and (c) displays the average word count in dataset documentation.

these variations reflect distinct cultural, institutional, and legal factors influencing dataset license norms across CJK and English communities.

Documentation Quality Patterns Figure 6 illustrates three key metrics related to dataset documentation. For academic grounding, measured by the presence of associated arXiv papers, English datasets lead with 39.20%, followed by Chinese (16.97%), while Japanese (10.48%) and Korean (8.89%) trail behind. Regarding structural completeness, measured by the presence of all five recommended Hugging Face dataset card sections, English again leads (18.41%), with Chinese (7.67%), Japanese (1.64%), and Korean (2.46%) showing lower completeness. In terms of documentation depth, measured by average word count, English maintains the highest average (678 words), Chinese stands at 354, and Japanese and Korean remain lower at 168 and 182 respectively. Taken together, these figures indicate a consistent trend: English datasets demonstrate the most thorough and standardized documentation practices, Chinese resources show moderate completeness and Japanese and Korean documentation remains comparatively succinct or under-documented.

Dataset Ownership Patterns Figure 7 illustrates the proportions of datasets contributed by individual and community contributors, research institutes, and companies in the Korean, Chinese, and Japanese communities. Individual and community contributors dominate across all three languages: Korean datasets lead with 79.2%, followed by Japanese at 71.9% and Chinese at 61.6%. Research institutes play a stronger role in China (27.3%)

HuggingFace dataset card sections: *Description, Structure, Creation, Usage, and Additional Info*
 Analysis limited to CJK datasets with verifiable institutional affiliations.

than in Korea (11.1%) or Japan (21.0%), suggesting more prominent institutional involvement in Chinese NLP resource development. Company contributions remain the smallest category across all three, though slightly higher in Chinese (11.1%) compared to Korean (9.7%) and Japanese (7.2%). This highlights the key role of grassroots efforts as primary drivers of dataset creation.

4.3 Language-Specific Characteristics

4.3.1 Chinese Dataset Ecosystem

Comprehensive Evaluation Frameworks Chinese NLP resources on Hugging Face frequently feature large-scale, well-structured evaluation suites. For example, *ceval/ceval-exam* (Huang et al., 2023) provides 13,948 multiple-choice questions across 52 domains, and *haonan-li/cmmlu* (Li et al., 2023b) covers 67 subject areas spanning elementary to advanced professional levels. These broad assessments facilitate detailed benchmarking of model performance across diverse knowledge domains.

Specialized Medical Domain Resources Chinese datasets also demonstrate significant depth in specialized fields. The *FreedomIntelligence/CM B* (Wang et al., 2023) collection features a hierarchical structure (6 main categories, 28 subcategories) spanning 11,200 medical questions, thus enabling targeted evaluations in clinical question-answering. Similarly, *TCMLM/TCM_Humanities* (Kang, 2024) focuses on Traditional Chinese Medicine, integrating resources from professional certification materials and historical texts.

Dialectal and Cultural Diversity In addition to specialized domains, Chinese datasets often emphasize linguistic diversity and cultural preservation. The *Nexdata/chinese_dialect* (Nexdata, 2025) corpus contains 25,000 hours of dialect

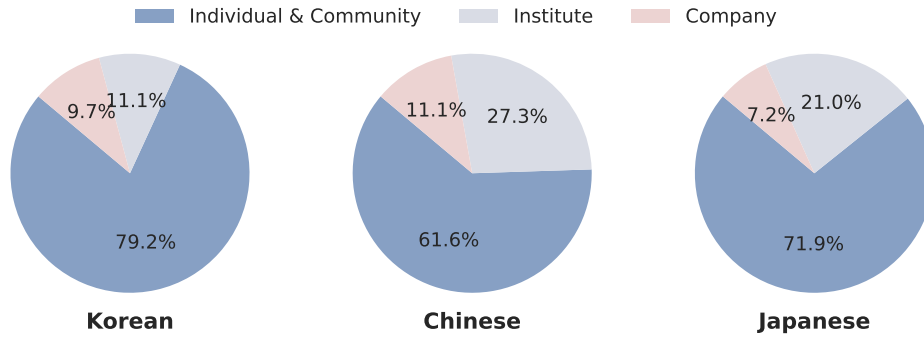


Figure 7: Dataset ownership across Korean, Chinese, and Japanese datasets, illustrating the proportions of individual and community contributors, research institutes, and companies.

speech data, facilitating fine-grained dialect modeling. Likewise, `raptorkwok/cantonese-traditional-chinese-parallel-corpus` ([raptorkwok, 2025](#)) offers over 130k aligned sentence pairs for Cantonese–Mandarin translation. Future efforts could enrich such dialectal resources by detailing speaker demographics, annotation workflows, and language-specific quirks—thereby promoting more equitable research coverage across China’s diverse linguistic communities.

4.3.2 Korean Dataset Ecosystem

Community-Driven Development and Its Impact Korean datasets on Hugging Face frequently emerge from grassroots, community-led efforts, rather than purely institutional or corporate projects. Prominent examples include contributions from open-source communities like HAERAE, which developed the widely-used HAERAE-HUB/KMMLU benchmark ([Son et al., 2024](#)), as well as individual developers such as `beomi`, `maywell`, and `taeminlee`, who have created highly-downloaded resources. Even widely-used benchmarks like `klue/klue` ([Park et al., 2021](#)) represent collaborative efforts between academia, industry, and individual researchers rather than single-entity projects. The broader ecosystem is dominated by individual and community contributors who account for 79.2% of Korean datasets. This community-driven approach has accelerated the proliferation of new resources but also contributed to inconsistencies in documentation and licensing. For instance, Korean has the highest proportion of “Unknown” licenses among CJK languages (35.3%), indicating gaps in legal clarity and potential challenges for commercial or cross-institutional usage. Moreover, only 8.89% of Korean datasets are linked to an arXiv publication—lower than both Chinese (16.97%)

and Japanese (10.48%). These factors may hinder collaborative research or reproducibility, underscoring the need for more standardized dataset cards ([Gebru et al., 2021](#)) and explicit licensing.

Content Moderation Focus A unique strength of the Korean dataset ecosystem is its emphasis on content moderation, encompassing hate-speech detection, toxicity filtering, and profanity masking. Popular resources such as `jeanlee/kmhas-korean-hate-speech` ([Lee et al., 2022](#)), `humane-lab/K-HATERS` ([Lab, 2025](#)) and `Dasool/KoMultiText` ([Choi et al., 2023](#)) reflect heightened community and research interest in combating harmful or discriminatory language online. However, these moderation-oriented resources raise broader ethical and regulatory questions, such as defining thresholds for *hate speech* or handling user privacy. Although the Hugging Face platform provides general community guidelines, more detailed policies—particularly for age-restricted or sensitive data—would help standardize safe usage of these potentially sensitive resources.

4.3.3 Japanese Dataset Ecosystem

Strong focus on subcultural content Japanese NLP datasets often integrate subcultural or entertainment-related material, an approach that distinguishes them from other CJK resources. For instance, `joujiboi/japanese-anime-speech` ([joujiboi, 2024](#)) targets automatic speech recognition in anime content, attracting high download counts and demonstrating direct utility for real-world applications such as subtitle generation. Additionally, `YANS-official/ogiri-test-with-references` ([YANS-official, 2023](#)) captures the distinctive *Ogiri* comedy tradition, illustrating

As of January 2025, based on Hugging Face’s publicly available platform policies and community guidelines.

Japan’s unique comedy culture through multimodal data (text and images). While such resources enrich models’ ability to handle colloquial or creative contexts, they also require careful documentation of stylistic nuances and potential copyright constraints. Many subcultural datasets involve fan works or licensed content, which often preclude fully open licenses. Researchers must therefore verify these constraints to avoid unintended usage restrictions or downstream complications.

Diverse Methods in Dataset Processing and Refinement Japanese datasets exhibit a reliance on translation-based pipelines and synthetic data generation rather than building new corpora from scratch. For example, the Magpie series (Xu et al., 2024) has been adapted into multiple Japanese resources—e.g., Aratako/Synthetic-JP-EN-Translation-Dataset-Magpie-Nemotron-4-20k and Aratako/Magpie-Tanuki-8B-annotated-96k (Aratako, 2024b,a)—highlighting how translations and AI-generated text can expand training data. While these strategies improve dataset availability, they raise concerns about translation errors, cultural misalignment, and potential biases introduced by synthetic text. Efforts such as neody/oscar-ja-cleaned (neody, 2023) and saillabalpaca-japanese-cleaned (Upadhayay and Behzadan, 2024) illustrate attempts to mitigate these issues through dataset cleaning and quality control. Systematic documentation of translation processes and validation protocols would help researchers assess dataset reliability. This localization approach may serve as a model for other languages seeking rapid resource expansion.

5 Discussion

Our analysis reveals distinct characteristics across CJK dataset ecosystems: Chinese datasets show strong institutional backing but inconsistent documentation; Korean datasets demonstrate community-driven development but face licensing gaps; and Japanese datasets emphasize subcultural content while dealing with copyright constraints.

Three practical issues warrant attention. First, licensing diversity (particularly "Unknown" licenses in Korean datasets and restricted licenses in Japanese resources) complicates collaborative projects. More consistent adherence to guidelines like "Datasheets for Datasets" (Geburu et al., 2021) could enhance reusability. Second, domain clustering (medical for Chinese, subcultural for Japanese,

and moderation for Korean) may underserve other areas needed for general-purpose LLM development. Third, culturally specific content requires transparent documentation, as simple translations miss nuanced cultural meanings.

Despite these differences, strong synergy potential exists. Joint benchmarks could facilitate cross-lingual comparisons, while unified documentation frameworks could standardize metadata and licensing. Our findings underscore both the richness and fragmentation of CJK resources, suggesting that clearer practices and cross-lingual collaboration can foster a robust ecosystem for East Asian LLM development.

6 Limitations

Our analysis primarily focused on datasets with relatively high download counts, which may have led us to overlook smaller or emerging resources that could shed light on niche trends or specialized applications. Furthermore, we limited our scope to the Hugging Face platform; investigating additional repositories (e.g., GitHub, Kaggle, or Papers with Code) could reveal a broader range of dataset characteristics and host factors. Although we manually requested permission to access certain private or restricted datasets, some ultimately remained inaccessible, thereby constraining the representativeness of our findings.

In addition, while Korean and Japanese datasets were examined with input from language experts, our review of Chinese data relied solely on documentation, potentially affecting the depth of our analysis. Finally, we chose to focus on three major East Asian languages, excluding many low-resource languages and dialects, whose inclusion could further expand and enrich our findings.

7 Conclusion

This study presents a comparative analysis of over 3,300 Chinese, Japanese, and Korean datasets on HuggingFace, revealing distinct ecosystem characteristics—Chinese datasets show strong institutional involvement, Korean resources are community-driven, and Japanese datasets emphasize subcultural content—highlighting that documentation, licensing, and ownership must be addressed in cultural context to guide inclusive East Asian language technologies.

Acknowledgements

This research was supported by Brian Impact, a non-profit organization dedicated to advancing science and technology.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Aratako. 2024a. Magpie-tanuki-8b-annotated-96k. <https://huggingface.co/datasets/Aratako/Magpie-Tanuki-8B-annotated-96k>. License: apache-2.0; Text Generation task; Japanese; 96.4k rows.
- Aratako. 2024b. Synthetic-jp-en-translation-dataset-magpie-nemotron-4-20k. <https://huggingface.co/datasets/Aratako/Synthetic-JP-EN-Translation-Dataset-Magpie-Nemotron-4-20k>. License: apache-2.0.
- Emily Bender. 2019. The# benderrule: On naming the languages we study and why it matters. *The Gradient*, 14:34.
- 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.
- Won Ik Cho, Sangwhan Moon, and Youngsook Song. 2020. Open korean corpora: A practical report. *arXiv preprint arXiv:2012.15621*.
- Won Ik Cho, Sangwhan Moon, and Youngsook Song. 2023. Revisiting korean corpus studies through technological advances. In *Proceedings of the 37th Pacific Asia Conference on Language, Information and Computation*, pages 653–660.
- Dasol Choi, Jooyoung Song, Eunsun Lee, Jinwoo Seo, Heejune Park, and Dongbin Na. 2023. Komultitext: Large-scale korean text dataset for classifying biased speech in real-world online services. *arXiv preprint arXiv:2310.04313*.
- Roberts Dargis, Guntis Barzdins, Inguna Skadiņa, and Baiba Saulīte. 2024. Evaluating open-source llms in low-resource languages: Insights from latvian high school exams. In *Proceedings of the 4th International Conference on Natural Language Processing for Digital Humanities*, pages 289–293.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, Junlei Zhang, Jinghan Zhang, Tangjun Su, Junteng Liu, Chuancheng Lv, Yikai Zhang, Jiayi Lei, Yao Fu, Maosong Sun, and Junxian He. 2023. C-eval: A multi-level multi-discipline chinese evaluation suite for foundation models. *arXiv preprint arXiv:2305.08322*.
- Hugging Face. 2023. [Hugging Face: The AI community building the future](#). Accessed: 2025-02-07.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The state and fate of linguistic diversity and inclusion in the nlp world. *arXiv preprint arXiv:2004.09095*.
- joujiboi. 2024. [japanese-anime-speech: An audio-text dataset for automatic speech recognition in anime](#). 73,004 audio-text pairs, 110 hours of audio. Licensed under CC0-1.0. Latest version: V5 - March 22nd 2024.
- Paris Kang. 2024. [Chinese medical humanities dataset](#). Hugging Face Dataset Hub.
- Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of naacL-HLT*, volume 1. Minneapolis, Minnesota.
- HUMANE Lab. 2025. [K-HATERS: Korean Hate and Offensive Speech Dataset](#). <https://huggingface.co/datasets/humane-lab/K-HATERS>. Accessed: 2025-01-28.
- Teven Le Scao, Angela Fan, Christopher Akiki, Elie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2023. Bloom: A 176b-parameter open-access multilingual language model.
- Jean Lee, Taejun Lim, Heejun Lee, Bogeun Jo, Yangsok Kim, Heegeun Yoon, and Soyeon Caren Han. 2022. [K-MHaS: A multi-label hate speech detection dataset in Korean online news comment](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3530–3538, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Quentin Lhoest, Albert Villanova Del Moral, Yacine Jernite, Abhishek Thakur, Patrick Von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, et al. 2021. Datasets: A community library for natural language processing. *arXiv preprint arXiv:2109.02846*.
- Anran Li, Weidong Zhan, Jia-Fei Hong, Zhao-Ming Gao, and Chu-Ren Huang. 2023a. Chinese language

- resources: A comprehensive compendium. In *Chinese Language Resources: Data Collection, Linguistic Analysis, Annotation and Language Processing*, pages 623–662. Springer.
- Haonan Li, Yixuan Zhang, Fajri Koto, Yifei Yang, Hai Zhao, Yeyun Gong, Nan Duan, and Timothy Baldwin. 2023b. [Cmmlu: Measuring massive multitask language understanding in chinese](#).
- neody. 2023. [oscar-ja-cleaned](https://huggingface.co/datasets/neody/oscar-ja-cleaned). <https://huggingface.co/datasets/neody/oscar-ja-cleaned>. A cleaned Japanese OSCAR dataset in Parquet format with 12.8M rows. Licensed under CC0-1.0.
- Nexdata. 2025. [chinese_dialect](https://huggingface.co/datasets/nexdata/chinese-dialect): A 25,000-hour chinese dialect speech dataset. <https://huggingface.co/datasets/nexdata/chinese-dialect>. Accessed: 2025-01-28.
- Sungjoon Park, Jihyung Moon, Sungdong Kim, Won Ik Cho, Jiyeon Han, Jangwon Park, Chisung Song, Junseong Kim, Yongsook Song, Taehwan Oh, et al. 2021. [Klue: Korean language understanding evaluation](#). *arXiv preprint arXiv:2105.09680*.
- Federica Pepe, Vittoria Nardone, Antonio Mastropaolo, Gabriele Bavota, Gerardo Canfora, and Massimiliano Di Penta. 2024. How do hugging face models document datasets, bias, and licenses? an empirical study. In *Proceedings of the 32nd IEEE/ACM International Conference on Program Comprehension*, pages 370–381.
- raptorkwok. 2025. [Cantonese-Traditional Chinese Parallel Corpus](https://huggingface.co/datasets/raptorkwok/cantonese-traditional-chinese-parallel-corpus). <https://huggingface.co/datasets/raptorkwok/cantonese-traditional-chinese-parallel-corpus>. Accessed: 2025-01-28.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. 2022. [Laion-5b: An open large-scale dataset for training next generation image-text models](#). *Advances in Neural Information Processing Systems*, 35:25278–25294.
- Guijin Son, Hanwool Lee, Sungdong Kim, Seungone Kim, Niklas Muennighoff, Taekyoon Choi, Cheonbok Park, Kang Min Yoo, and Stella Biderman. 2024. [Kmmlu: Measuring massive multitask language understanding in korean](#). *arXiv preprint arXiv:2402.11548*.
- Jianhua Tao, Fang Zheng, Aijun Li, and Ya Li. 2009. [Advances in chinese natural language processing and language resources](#). In *2009 Oriental COCOSA International Conference on Speech Database and Assessments*, pages 13–18. IEEE.
- Bibek Upadhyay and Vahid Behzadan. 2024. [Taco: Enhancing cross-lingual transfer for low-resource languages in LLMs through translation-assisted chain-of-thought processes](#). In *5th Workshop on practical ML for limited/low resource settings, ICLR*.
- Xidong Wang, Guiming Hardy Chen, Dingjie Song, Zhiyi Zhang, Qingying Xiao, Xiangbo Wu, Feng Jiang, Jianquan Li, and Benyou Wang. 2023. [Cmb: Chinese medical benchmark](https://github.com/FreedomIntelligence/CMB). <https://github.com/FreedomIntelligence/CMB>. Xidong Wang, Guiming Hardy Chen, Dingjie Song, and Zhiyi Zhang contributed equally.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Edouard Grave. 2019. [Ccnet: Extracting high quality monolingual datasets from web crawl data](#). *arXiv preprint arXiv:1911.00359*.
- Zhangchen Xu, Fengqing Jiang, Luyao Niu, Yuntian Deng, Radha Poovendran, Yejin Choi, and Bill Yuchen Lin. 2024. [Magpie: Alignment data synthesis from scratch by prompting aligned llms with nothing](#).
- Xinyu Yang, Weixin Liang, and James Zou. 2024. [Navigating dataset documentations in ai: A large-scale analysis of dataset cards on hugging face](#). *arXiv preprint arXiv:2401.13822*.
- YANS-official. 2023. [ogiri-test-with-references: Bokete crawl data](#). This dataset is crawled from the Bokete website and is a subset of the CLoT-Oogiri-Go [Zhang+ CVPR2024] dataset. It includes tasks for Image-to-Text and Text2Text Generation, with the test split consisting of 100 rows. Data preprocessing and filtering were applied as described in the dataset card.