

# CRAFT: Extracting and Tuning Cultural Instructions from the Wild

Bin Wang<sup>♡</sup>, Geyu Lin<sup>♡</sup>, Zhengyuan Liu<sup>♡</sup>, Chengwei Wei<sup>§</sup>, Nancy F. Chen<sup>♡,†</sup>

<sup>♡</sup>Institute for Infocomm Research (I<sup>2</sup>R), A\*STAR, Singapore

<sup>§</sup>University of Southern California, USA

<sup>†</sup>Centre for Frontier AI Research (CFAR), A\*STAR, Singapore

wang\_bin@i2r.a-star.edu.sg

## Abstract

Large language models (LLMs) have rapidly evolved as the foundation of various natural language processing (NLP) applications. Despite their wide use cases, their understanding of culturally-related concepts and reasoning remains limited. Meantime, there is a significant need to enhance these models' cultural reasoning capabilities, especially concerning underrepresented regions. This paper introduces a novel pipeline for extracting high-quality, culturally-related instruction tuning datasets from vast unstructured corpora. We utilize a self-instruction generation pipeline to identify cultural concepts and trigger instruction. By integrating with a general-purpose instruction tuning dataset, our model demonstrates enhanced capabilities in recognizing and understanding regional cultural nuances, thereby enhancing its reasoning capabilities. We conduct experiments across three regions: Singapore, the Philippines, and the United States, achieving performance improvement of up to 6%. Our research opens new avenues for extracting cultural instruction tuning sets directly from unstructured data, setting a precedent for future innovations in the field.<sup>1</sup>

## 1 Introduction

Large language models (LLMs) like ChatGPT (Achiam et al., 2023), Claude, and Gemini (Reid et al., 2024) have demonstrated their proficiency in managing diverse tasks related to semantic understanding and text generation. Beyond acting as general task-solvers, the ability of LLMs to understand and reason with cultural nuances could play a crucial role in generating precise and personalized responses to benefit broader communities (Tao et al., 2023; Wang et al., 2024; Adilazuarda et al., 2024).

Culture is a comprehensive concept encompassing traditions, customs, beliefs, values, and social

norms, all deeply rooted in historical contexts and continuously evolving over time. It is also intrinsically linked to languages and dialects, which can be sparsely represented in available resources.

In the domain of LLMs, which initially train on vast amounts of unlabeled data, knowledge is systematically captured and structured through data-driven techniques. With limited model sizes, knowledge that occurs infrequently is often less effectively captured compared to more frequently occurring information (Kaplan et al., 2020). Additionally, the predominance of English in the pre-training corpus inherently biases these models towards Western perspectives, a consequence of the over-representation of English-language sources. This bias means that cultural concepts may be inadequately captured, especially for under-represented regions (Masoud et al., 2023). Consequently, LLMs struggle to effectively adapt to and represent diverse cultural concepts due to these inherent limitations in their training data.

Expanding the cultural reasoning capabilities of LLMs could potentially be achieved by pre-training them on corpora from diverse languages. However, this approach is still expansive and challenging due to the difficulty in obtaining high-quality multilingual datasets (Bai et al., 2023; Singapore, 2023). Meanwhile, instruction fine-tuning could more directly impact end-user applications. However, the development of cultural instruction tuning sets is limited due to the high costs associated with collecting culturally relevant instruction sets, along with challenges in ensuring their quality and diversity.

In this study, we pioneer the study of deriving instruction tuning sets from unlabeled corpora. Initially, we use keyword filtering to isolate culturally relevant concepts from a vast corpus containing over 600 billion English tokens. Subsequently, we then utilized these selected regional text segments to prompt LLMs for both questions and answers. Our evaluations focus on the context of Singapore

<sup>1</sup>Our models and datasets are available for future research at <https://github.com/SeaEval/CRAFT>.

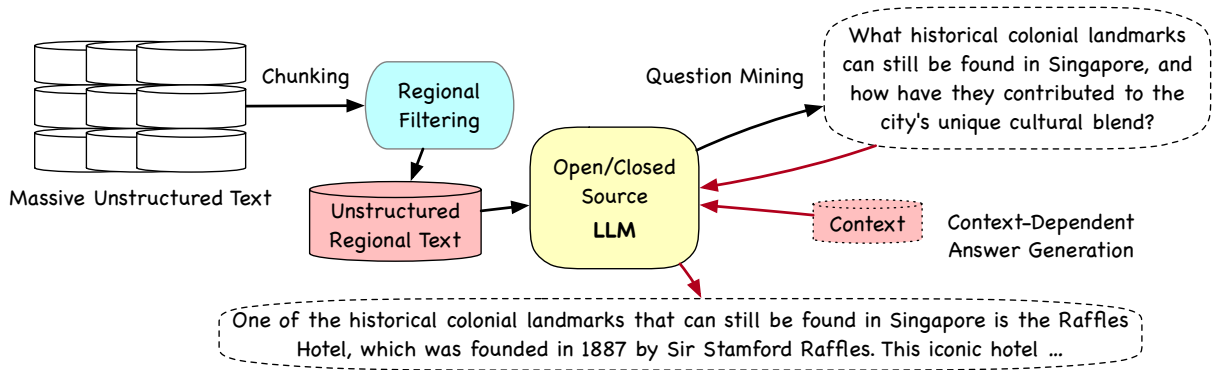


Figure 1: The CRAFT method involves creating instruction datasets tailored for culturally rich instruction by processing extensive unstructured data with large language models (LLMs). These specialized cultural instructions are then employed to improve the ability of LLMs to reason within cultural contexts through instruction fine-tuning.

and extend to the Philippines and the US. Our experiments utilize the *Mistral-7B* model, combining general instructions with our specifically curated cultural instruction set. We observe performance improvement of up to 6% while maintaining intelligence in general subject knowledge as assessed by the MMLU dataset. Additionally, we also analyzed the impact of answer sources and the ratio of cultural instructions. Both the model and the curated cultural instruction-tuning dataset are made available for future research.

## 2 Related Work

Our work aims to improve the cultural reasoning capability of LLMs. The capability of LLMs can be refined in their training schemes, such as pre-training, instruction tuning, and RLHF preference optimization. Current efforts often concentrate on fine-tuning using specific datasets derived from diverse multilingual corpus sources (Lin and Chen, 2023; Abbasi et al., 2023; Pipatanakul et al., 2023). Yet, this approach is costly in training and lacks transparency regarding the extent of cultural concepts incorporated into the model. Li et al. (2024) propose to leverage a set of opinion questions to gather views towards different cultural groups, which is then used to finetune a model. However, despite the efforts at data augmentation, the scope of these questions remains limited and fails to encompass a wide range of cultural dimensions. Therefore, in our research, we focus on extracting a wide variety of cultural instructions from large unlabeled corpora, ensuring guaranteed diversity.

Given the complexity of sourcing high-quality instructional data, LLMs are employed to create

synthetic question-answering pairs (Wang et al., 2022, 2023) and dialogue data through iterative processes (Ding et al., 2023). However, these approaches tend to concentrate on generating general instructions and dialogues, lacking the capability to produce culturally rich instructions. Prompting LLMs to directly generate cultural concepts is also challenging, as these concepts are sparsely distributed across various resources.

## 3 Methodology

We introduce the **CRAFT** (Cultural Reasoning with Instruction Fine-Tuning) method, designed to synthesize cultural instructions from a massive, unlabeled English corpus. The methodology is detailed in the following steps, as illustrated in Figure 1.

**Selective Data Extraction.** We utilize SlimPajama (Soboleva et al., 2023) as our primary data source, which comprises an English corpus containing over 600 billion tokens. Given the sparse distribution of culturally relevant concepts within this vast dataset, and to manage the processing burden effectively, we propose an efficient filtering process using keywords to identify and extract culture-related concepts.

Specifically, we curate a collection comprising a minimum of 150 words to represent each region. We then segment the documents into chunks no larger than 512 tokens each. From these segments, we retain only those text chunks that include at least two regional keywords, such as "National Day Parade" and "Merlion" for Singapore. Through analyzing a subset of over 200 billion English tokens, we successfully extracted 35,000 text segments for Singapore, 25,000 for the Philippines, and 35,000

Models	SG-Eval	PH-Eval	US-Eval	MMLU
<b>General LLMs</b>				
<i>ChatGPT-3.5</i>	64.4	58.4	74.9	67.5
<i>LLaMA-3-8B-Instruct</i> (Meta_AI, 2024)	62.1	54.6	69.8	62.5
<i>LLaMA-2-7B-Chat</i> (Touvron et al., 2023)	39.8	35.4	50.1	44.5
<i>Mistral-7B-Instruct-v0.2</i>	62.1	48.6	60.9	58.1
<b>Base Model: <i>Mistral-7B-Instruct-v0.2</i> (Jiang et al., 2023)</b>				
Tuning w/ OpenHermes-2.5	62.6	46.4	65.5	58.5
CRAFT <sub>sg</sub>	<b>68.3</b> / 64.2	46.2 / 44.8	65.2 / 64.6	59.8 / 58.4
CRAFT <sub>ph</sub>	64.3 / 63.7	<b>49.2</b> / 44.6	65.4 / 64.7	60.0 / 59.4
CRAFT <sub>us</sub>	65.3 / 63.2	48.4 / 44.8	<b>67.1</b> / 63.5	60.3 / 59.7

Table 1: The main results for general LLMs and our model across three cultural evaluation datasets and the MMLU dataset, which assesses general knowledge. For our results, "-/-" denotes the scores for context-dependent answers and context-free answers, respectively.

for the US.

**Automated Question Creation.** Given the text chunks rich in cultural and local content, we prompt an off-the-shelf LLM to generate questions specifically related to each chunk, focusing on the cultural and regional concepts mentioned.

**Answer Production.** To collect responses for the generated questions, we employ two approaches: 1) context-dependent answer generation, where the given context is provided to LLMs when forming answers, as shown in Figure 1, and 2) context-free answer generation, allowing for responses that are more creative and less tailored to the immediate context. For automatic question creation and context-dependent answer generation, we utilize *Zephyr-7B-Beta* (Tunstall et al., 2023). For context-free answer generation, we employ *ChatGPT-3.5*.

**Hybrid Instruction Tuning.** After developing the cultural instructions, we compiled at least 20,000 instructions for each specified region. To ensure a balanced capability, we incorporated 50,000 single-round instructions from the OpenHermes-2.5 (Teknum, 2023) dataset alongside random sampled 20,000 cultural instructions to fine-tune the *Mistral-7B-Instruct-v0.2* (Jiang et al., 2023) model. We adapt LoRA (Hu et al., 2021) training with a rank of 16 to effectively integrate our instruction tuning knowledge into the model.

## 4 Experiments

**Settings.** For evaluation, we selected four datasets, including three culturally-focused ones: SG-Eval, Ph-Eval, and US-Eval (Wang et al., 2024). Each

of these datasets is a human-crafted collection of multiple-choice questions and answers that probe cultural or regional knowledge. To assess knowledge capabilities in general subjects, we utilize a subset of the MMLU dataset (Hendrycks et al., 2020; Wang et al., 2024). To minimize the influence of randomness, we report the average results from five runs, using paraphrased prompt templates.

**Results and Discussion.** Table 1 shows the main result where we listed four competitive LLMs and our baseline model tuned with 50,000 OpenHermes instructions. First, we witness that *ChatGPT-3.5* and recently released *LLaMA-3-8B-Instruct* are strong in capturing cultural concepts from different regions. Compared to *Mistral-7B* models, LLaMA-3’s larger training data and multilingual tokens could help improve the model’s capability.

Compared to our baseline, our CRAFT model demonstrates a clear performance improvement across all three regions. Notably, we observe more significant enhancements for Singapore-related questions, improving from 62.6% to 68.3%. We attribute this to two main reasons: 1) It is a culturally rich region with mixed and unique cultural origins, and 2) Singaporean concepts are well-documented in English corpora, even though they constitute a small portion. In contrast, cultural knowledge of the US appears more frequently during the pretraining stage, leading to better learning outcomes. Conversely, cultural knowledge from the Philippines is less documented and more dispersed. We encountered similar challenges when collecting culturally related text chunks for it. We found that identifying

Region	Instruction and Answer
SG	<b>Q:</b> What historical significance does the Mandarin Orchard Singapore hold, and how has it been preserved while undergoing a major refurbishment? <b>A1:</b> The Mandarin Orchard Singapore holds historical significance as it is a landmark building on Orchard Road, with its two distinctive towers built in 1971 and 1973. ... <b>A2:</b> The Mandarin Orchard Singapore, previously known as the Mandarin Singapore, holds significant historical and cultural importance as one of Singapore’s ...
SG	<b>Q:</b> How has Kuala Lumpur’s mix of Malay, Indian, and Chinese cultures contributed to its unique experience and creative options for event planners, as mentioned in the context about Kuala Lumpur being a city of vibrant contrasts? <b>A2:</b> Kuala Lumpur’s blend of Malay, ...
PH	<b>Q:</b> How is the Araquo festival celebrated in Nueva Ecija, and what is its traditional month of celebration? <b>A1:</b> The Araquo festival is traditionally celebrated in the month of May in Nueva ...
US	<b>Q:</b> How did the discovery of gold in California in 1848 impact the transition from Mexican ranching to American farms and towns in the East Bay area, and what were the consequences for the Estudillo and Peralta families who owned land in the region? <b>A2:</b> The discovery of gold in California ...

Table 2: Synthesized cultural instructions and answers. **A1** and **A2** refer to context-dependent and context-free answers, respectively.




Figure 2: Performance on SG-Eval and MMLU dataset. The CRAFT method with different ratios of Singapore cultural instruction data.

US text chunks required the least data, followed by Singapore, with the Philippines proving to be the most challenging due to its scattered documentation.

Lastly, when comparing context-dependent answers with context-free answers, we found that responses derived from context are consistently more reliable and informative. Consequently, instructions that utilize context-dependent answers consistently yield higher performance gains. While context-free answers could potentially become more informative using advanced models like GPT-4, this enhancement would also benefit context-dependent answers.

**Cultural Instruction Ratios.** In Figure 2, we investigate the effects of adding varying amounts of cultural instructions to a base of 50,000 general

instructions. Cultural instructions are incrementally introduced into the training data at intervals of 2,500 samples. The results indicate that performance on the SG-Eval improves as more culturally related samples are added, suggesting that an increased number of cultural concepts are activated from the pre-training phase and further improved from the instructions. Concurrently, our analysis shows that the general performance as measured by the MMLU datasets remains consistent.

**Instruction Samples.** In this section, we present a series of synthesized cultural instructions and their corresponding responses generated by various methods, as illustrated in Table 2. The samples demonstrate that culturally related questions can be effectively derived from culturally specific content. Context-dependent responses tend to incorporate more factual knowledge from the provided context compared to context-free answers. However, this can also lead to biases based on the limited facts presented. We observe that context-free responses often lack depth in knowledge-intensive instructions due to the model’s inherent limitations.

## 5 Conclusion

In this paper, we introduce the CRAFT method, designed to synthesize cultural instructions from a vast, unlabeled corpus. We conduct experiments across three regions, with the potential for expansion to additional regions. Our pioneering self-instruction techniques facilitate effective mining from unstructured data sources, enhancing both the diversity and quality of the synthesized instructions compared to previous studies.

## Limitations

In this study, we concentrate on mining cultural instructions from English-language corpora. However, it is important to recognize that cultural concepts are often deeply integrated with their respective languages, including those that are primarily spoken. Therefore, to effectively synthesize cultural concepts and instructions, adopting multilingual approaches (Liu et al., 2023; Lin et al., 2024) is essential to accommodate a broader range of cultural contexts.

## Acknowledgement

This work is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-GC-2022-005). This research is also supported by the National Research Foundation, Singapore and Infocomm Media Development Authority, Singapore under its National Large Language Models Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore and Infocomm Media Development Authority, Singapore.

## References

- Mohammad Amin Abbasi, Arash Ghafouri, Mahdi Firouzmandi, Hassan Naderi, and Behrouz Minaei Bidgoli. 2023. Persianllama: Towards building first persian large language model. *arXiv preprint arXiv:2312.15713*.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Muhammad Farid Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Singh, Ashutosh Dwivedi, Alham Fikri Aji, Jacki O’Neill, Ashutosh Modi, and Monojit Choudhury. 2024. Towards measuring and modeling" culture" in llms: A survey. *arXiv preprint arXiv:2403.15412*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Shengding Hu, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. 2023. Enhancing chat language models by scaling high-quality instructional conversations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3029–3051, Singapore. Association for Computational Linguistics.
- 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*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- 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*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Cheng Li, Mengzhou Chen, Jindong Wang, Sunayana Sitaram, and Xing Xie. 2024. Culturellm: Incorporating cultural differences into large language models. *arXiv preprint arXiv:2402.10946*.
- Geyu Lin, Bin Wang, Zhengyuan Liu, and Nancy F Chen. 2024. Crossin: An efficient instruction tuning approach for cross-lingual knowledge alignment. *arXiv preprint arXiv:2404.11932*.
- Yen-Ting Lin and Yun-Nung Chen. 2023. Taiwan llm: Bridging the linguistic divide with a culturally aligned language model. *arXiv preprint arXiv:2311.17487*.
- Chen Cecilia Liu, Fajri Koto, Timothy Baldwin, and Iryna Gurevych. 2023. Are multilingual llms culturally-diverse reasoners? an investigation into multicultural proverbs and sayings. *arXiv preprint arXiv:2309.08591*.
- Reem I Masoud, Ziquan Liu, Martin Ferianc, Philip Treleaven, and Miguel Rodrigues. 2023. Cultural alignment in large language models: An explanatory analysis based on hofstede’s cultural dimensions. *arXiv preprint arXiv:2309.12342*.
- Meta\_AI. 2024. Meta llama 3. <https://ai.meta.com/blog/meta-llama-3/>. Accessed: 2024-05-01.
- Kunat Pipatanakul, Phatrasek Jirabovonvisut, Potsawee Manakul, Sittipong Sripaisarnmongkol, Ruangsak Patomwong, Pathomporn Chokchainant, and Kasima Tharnpipitchai. 2023. Typhoon: Thai large language models. *arXiv preprint arXiv:2312.13951*.

Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.

AI Singapore. 2023. Sea-lion (southeast asian languages in one network): A family of large language models for southeast asia. <https://github.com/aisingapore/sealion>.

Daria Soboleva, Faisal Al-Khateeb, Robert Myers, Jacob R Steeves, Joel Hestness, and Nolan Dey. 2023. SlimPajama: A 627B token cleaned and deduplicated version of RedPajama.

Yan Tao, Olga Viberg, Ryan S Baker, and Rene F Kizilcec. 2023. Auditing and mitigating cultural bias in llms. *arXiv preprint arXiv:2311.14096*.

Teknum. 2023. Openhermes 2.5: An open dataset of synthetic data for generalist llm assistants.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrubti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra, Cl  mentine Fourrier, Nathan Habib, et al. 2023. Zephyr: Direct distillation of lm alignment. *arXiv preprint arXiv:2310.16944*.

Bin Wang, Zhengyuan Liu, and Nancy Chen. 2023. Instructive dialogue summarization with query aggregations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7630–7653, Singapore. Association for Computational Linguistics.

Bin Wang, Zhengyuan Liu, Xin Huang, Fangkai Jiao, Yang Ding, Ai Ti Aw, and Nancy F Chen. 2024. Seaval for multilingual foundation models: From cross-lingual alignment to cultural reasoning. *NAACL*.

Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hananeh Hajishirzi. 2022. Self-instruct: Aligning language models with self-generated instructions. *arXiv preprint arXiv:2212.10560*.

## A Keywords and Templates

This section details the keywords and instructions employed for selective data extraction, automated question generation, and answer formulation.

Regarding cultural keywords, examples are provided below. Our work focuses on the pioneering effort to synthesize cultural instructions. Thus,

these keywords can be further refined to enhance performance.

For keywords, a complete list can be found in our source code.

- **Singapore:** "HDB flats", "Bukit Merah", "Jurong West", "Orchard Road", "Marina Bay Sands", "Merlion", "Sentosa Island", "CPF (Central Provident Fund)", "Temasek", "Lee Kuan Yew", "BTO (Build-To-Order)", "SCDF (Singapore Civil Defence Force)", "Majulah Singapura (National Anthem)", etc.
- **the Philippines:** "Jeepney", "Barong Tagalog", "Sinulog Festival", "Adobo", "Balut", "OFW (Overseas Filipino Worker)", "Bayanihan", "Boracay", "Ifugao Rice Terraces", "Bahay Kubo", "Tinikling", etc.
- **USA:** "The Statue of Liberty", "Hollywood", "Route 66", "The Grand Canyon", "Mount Rushmore", "The Civil Rights Movement", "Mardi Gras", "The White House", "NASCAR", "Manhattan Project", "Broadway", "Apollo Moon Landing", "Civil War", "Harlem Renaissance", etc.

We employ a straightforward prompt for question generation: "You are a chatbot who always generates just one question about Singapore from the given context. Do not generate the answer.". For generating both context-dependent and context-free answers, we use the direct prompt: "Please answer the following question.". The difference lies in whether the context is provided and the model will decide on how much the answer should rely on the provided context.