



MAKIEVAL: A Multilingual Automatic Wikidata-based Framework for Cultural Awareness Evaluation for LLMs

Raoyuan Zhao, Beiduo Chen, Barbara Plank, Michael A. Hedderich

MaiNLP, Center for Information and Language Processing, LMU Munich
Munich Center for Machine Learning (MCML)
{rzhao, beiduo chen, bplank, hedderich}@cis.lmu.de

Abstract

Large language models (LLMs) are used globally across many languages, but their English-centric pretraining raises concerns about cross-lingual disparities for cultural awareness, often resulting in biased outputs. However, comprehensive multilingual evaluation remains challenging due to limited benchmarks and questionable translation quality. To better assess these disparities, we introduce MAKIEVAL, an automatic multilingual framework for evaluating cultural awareness in LLMs across languages, regions, and topics. MAKIEVAL evaluates open-ended text generation, capturing how models express culturally grounded knowledge in natural language. Leveraging Wikidata’s multilingual structure as a cross-lingual anchor, it automatically identifies cultural entities in model outputs and links them to structured knowledge, enabling scalable, language-agnostic evaluation without manual annotation or translation. We then introduce four metrics that capture complementary dimensions of cultural awareness: granularity, diversity, cultural specificity, and consensus across languages. We assess 7 LLMs developed from different parts of the world, encompassing both open-source and proprietary systems, across 13 languages, 19 countries and regions, and 6 culturally salient topics (e.g., food, clothing). Notably, we find that models tend to exhibit stronger cultural awareness in English, suggesting that English prompts more effectively activate culturally grounded knowledge.

1 Introduction

The global deployment of LLMs in applications ranging from conversational agents to content creation makes their cultural awareness increasingly critical (Pawar et al., 2024). Without such awareness, models risk producing biased or culturally insensitive outputs, which can perpetuate stereotypes and harm user trust (Bender et al., 2021; Hutchinson et al., 2020; Blodgett et al., 2020). Yet despite

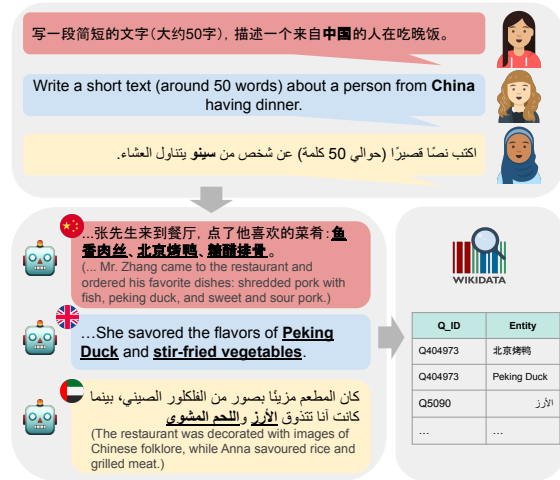


Figure 1: Multilingual prompts and corresponding model outputs for a culturally grounded generation task. Our culture-aware evaluation focuses on searching and assessing entities within these outputs using Wikidata.

its importance, we ask fundamentally: *How can we evaluate a text generation model’s cultural awareness in a flexible and generalizable manner?*

Figure 1 presents a realistic scenario: when prompted in different languages to describe a Chinese person having dinner, the same LLM produces outputs that differ in food entity specificity (i.e., whether the food items are culturally specific to China), diversity, and cross-lingual consistency. Capturing these differences across languages in evaluation, especially under the flexible nature of text generation models, remains a major challenge. Particularly, we identify three key limitations:

i) *test format*: dominant evaluation methods often rely on simplified setups such as cloze tests, single-token prediction, or predefined cultural entity lists (Naous et al., 2024; Wang et al., 2024). While computationally efficient, these formats fail to capture how cultural knowledge naturally manifests in language generation (Jones and Galliers, 1995; Bhatt and Diaz, 2024).

ii) *language and cultural coverage*: most exist-

ing studies focus on English, adopt monolingual setups (Onohara et al., 2025; Seveso et al., 2025; Naous et al., 2024; Nikandrou et al., 2025), or depend on translation-based pipelines for multilingual evaluation (Berger and Ponti, 2025). Although translation offers convenience, it introduces errors or meaning loss—particularly with culturally specific terms (Yao et al., 2024).

iii) *evaluation benchmark*: many existing benchmarks rely on manually curated datasets with fixed cultural entities (Schneider and Sitaram, 2024; Nayak et al., 2024). This static design lacks scalability and fails to adapt to emerging or less commonly documented cultural expressions, leading to incomplete assessments.

To address these limitations, we introduce MAKIEVAL, an automatic approach for i) open-ended text evaluation, ii) leveraging Wikidata as a language bridge and iii) enabling structured evaluation metrics and scalable, up-to-date evaluation. Instead of static benchmarking, MAKIEVAL is inspired from a human perspective and prior work: when humans describe daily life in their own region, they tend to be more detailed (high granularity), mention a wider range of items (high diversity), refer to culturally specific concepts (high culture specificity), and maintain consistent representations across languages (high cultural consensus) (Chafe, 1994; Dewaele and Pavlenko, 2003; Peters et al., 2018). For example, when asked about Chinese breakfast, an English speaker unfamiliar with the culture might respond simply “A Chinese eats noodles in the morning,” whereas a Chinese speaker would provide a richer description such as “youtiao (fried dough sticks) with soy milk, jianbing, or millet porridge in winter.” This contrast highlights all four dimensions. These dimensions have long been discussed in psycholinguistics and linguistics, and they also resonate with recent NLP research on semantic abstraction, coverage, cultural awareness, and cross-lingual consistency. In what follows, we draw on this line of work to design four operational metrics that capture cultural awareness in multilingual generation.

Building on these metrics, we conduct comprehensive experiments to investigate how model origin, prompt language, and the mentioned country affect LLMs’ cultural awareness. We found models align more with culturally related languages and their regions of origin. We publicly release our

code and data.¹ Our contributions are:

- We propose MAKIEVAL, an automatic framework for evaluating cultural awareness in multilingual text generation, accompanied by four distinct cultural-aware metrics: granularity, diversity, culture specificity, and culture consensus.
- We enable direct cross-linguistic comparison without relying on translation and conduct a large-scale evaluation in 13 languages across 19 countries and 7 models to reveal patterns of cultural awareness.
- We observe differences in cultural awareness may arise from a mix of factors, including model-, country- and language-specific characteristics.

2 Related Work

Dimensions of Cultural Evaluation While “culture” has become a prominent focus in natural language processing, existing work evaluates cultural awareness from widely varying dimensions. Liu et al. (2024b) address this fragmentation by proposing a taxonomy of cultural elements, organized into ideational, linguistic, and social categories, spanning both concrete and abstract phenomena.

Some investigations examine abstract or high-level aspects of culture, such as value systems (Ma et al., 2024; Agarwal et al., 2024; Shi et al., 2024), political ideology (Bang et al., 2024), gender norms (Wan et al., 2023), or expressions of stereotypes or hate (Cheng et al., 2023; Deshpande et al., 2023). Others focus on concrete, content-level evaluations, assessing how LLMs engage with specific cultural knowledge or practices (Naous et al., 2024; Wang et al., 2024; Jiang and Joshi, 2024; Zhou et al., 2025; Hu et al., 2024; Ghaboura et al., 2025). In this work, we also focus on the ideational branch, particularly on concrete cultural artifacts such as food, music, and clothing, which allow for a more structured, multilingual evaluation.

Benchmarks for Cultural Awareness And Bias

To systematically evaluate cultural bias and awareness in LLMs, researchers have proposed a range of benchmarks by pre-collecting data and constructing dedicated benchmarks (Yüksel et al., 2024; Jin et al., 2024; Chiu et al., 2024; Singh et al., 2025). Here, we highlight several recent works. JM-MMU (Onohara et al., 2025) focuses on Japanese culture awareness. ITALIC (Seveso et al., 2025)

¹<https://github.com/mainlp/MAKIEval>

targets cultural and linguistic understanding in Italian through a large-scale multiple-choice benchmark. CROPE (Nikandrou et al., 2025) constructs a culturally grounded multimodal benchmark by collecting concepts from Wikipedia and generating 1,060 binary-choice English questions that ask whether an image corresponds to a culturally associated entity. Naous et al. (2024) analyze cultural bias in Arabic prompts by examining how language models balance Arab and Western representations across three tasks: stereotype generation (via adjective sentiment), fairness in NER and sentiment classification, and cultural entity prediction in masked tweets. Predicted entities are compared with the CAMEL Cultural Entities dataset, where overuse of Western entities signals bias.

However, these benchmarks are mostly single-language, multiple-choice in format, and often cover only one specific country or region. Their reliance on pre-defined questions also limits generalization to open-ended settings.

Cross-Lingual Bias Comparison Some recent works extend the evaluation to multilingual settings. Myung et al. (2024) introduced a multilingual benchmark primarily consisting of short-answer and multiple-choice questions. They found that for resource-rich languages, LLMs tend to perform better in their native language. Berger and Ponti (2025) compared textual descriptions of the same images produced by human speakers of different languages, using tools like Google Translate and WordNet to facilitate alignment. Their analysis revealed cross-cultural differences in what aspects people focus on when describing the same visual content. Wang et al. (2024) prompted LLMs in different languages to list cultural artifacts like holidays, then identify the country attributes with the help of Wikipedia. Rather than constraining output formats to entity listing, our framework utilizes natural, open-ended prompts to guide the model in generating coherent scene descriptions. We presented the most relevant works to our setting here. For a more comprehensive overview, we refer to the survey by Pawar et al. (2025).

3 MAKIEVAL

We propose MAKIEVAL, an automatic and multilingual evaluation framework for evaluating cultural awareness. In Figure 2, we illustrate our pipeline, which consists of four steps: text generation, cultural entity extraction, wikidata-based entity match

and metric-based analysis.

3.1 Text Generation

We prompt LLMs to generate culturally grounded texts in multiple languages. Unlike prior evaluations that rely on constraint output formats such as cloze tasks or entity listing, our setup encourages open-ended generation, allowing models to produce more realistic outputs that reflect typical usage scenarios and cultural entities embedded in natural language text.

To enable systematic comparison across languages, topics, and countries/regions, we introduce a set of template structures for the input prompts. Each template contains a placeholder for a country or region name, which we instantiate dynamically to create culturally contextualized prompts. Importantly, the templates only serve to initialize the generation context—models are otherwise unconstrained in their output.

For each (language, topic, country) combination, we generate 500 responses per model to capture generation variance. Section 5.3 details the prompt design, including template examples and strategies for ensuring cross-lingual alignment.

3.2 Cultural Entity Extraction

We extract topic-relevant cultural entities from each generated text using GPT-4o-mini (OpenAI, 2025). For each topic (e.g., food, books, or music), GPT-4o-mini is prompted to identify mentions of relevant entities in the text (e.g., dish names for the food topic or author names for the book topic). In addition to extraction, the model classifies each entity by its semantic granularity. For example, for the food topic, it distinguishes whether the mention refers to a specific dish (e.g., Kung Pao Chicken), a dish type (e.g., stir-fry), an individual ingredient (e.g., onion), or an ingredient class (e.g., vegetable).

This process allows us to capture not only what cultural elements the model generates, but also how specific or abstract those references are—an essential signal for evaluating the depth and precision of cultural awareness.

3.3 Wikidata-Based Entity Match

To normalize entity mentions across languages and surface forms, we match each extracted entity to its corresponding entry in Wikidata (Vrandečić and Krötzsch, 2014). Our retrieval process matches not only entity labels but also aliases (“ALSO KNOWN

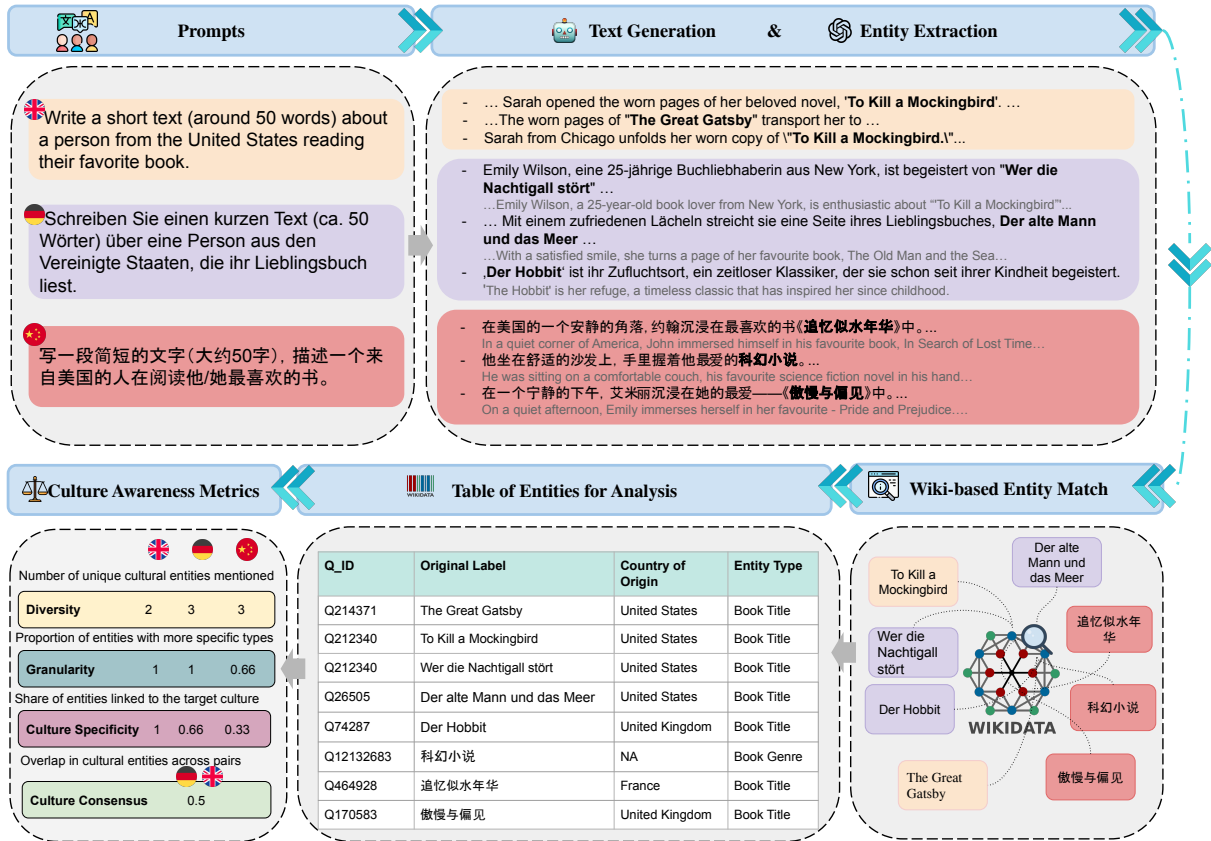


Figure 2: Overview of the MAKIEVAL, a multilingual, generation-based framework for evaluating cultural awareness in LLMs. Given prompts in different languages to LLM, MAKIEVAL collects open-ended generations, extracts named cultural entities and matched to Wikidata entries. Matched entities are augmented with metadata for downstream metric computation and analysis.

AS”) in the prompt’s language, falling back to English if no match is found. Each Wikidata entry is identified through its unique QID.

Since Wikidata contains ambiguous labels (e.g., “lemon” may refer to a fruit or a film, each with a different QID), we implement a disambiguation procedure that leverages hierarchical semantic properties. For each matched label, we collect all candidate QIDs and perform a recursive traversal over their instance of, subclass of, and part of relations. Each candidate may be linked to multiple such properties, resulting in a tree of semantic paths. We traverse all reachable paths upward in the ontology graph and retain the entity if any path leads to a culturally relevant category (e.g. food, literary work). If no valid semantic path is found, the disambiguation is marked as unresolved, and the corresponding entity is discarded from evaluation.

Once the correct entity is identified, we retrieve a set of metadata from Wikidata. This includes the entity’s description, country of origin, and, where applicable, the nationality of the author or per-

former (e.g., for books or music).

All matched entities and their corresponding metadata are compiled into a structured table that supports both quantitative and qualitative analysis. Each entry includes the entity’s QID, surface labels in multiple languages, granularity tags, and country/region-level information. A manual evaluation of the performance indicated strong results, with full details presented in Appendix C.

4 Evaluation Metrics

We propose four metrics to evaluate different aspects of cultural awareness in multilingual language models: **Granularity**, **Diversity**, **Culture Specificity**, and **Culture Consensus**. All metrics operate over a set of cultural entities extracted from the evaluated model’s responses.

We denote the union of all predicted entities across responses of a specific evaluated LLM and a specific prompt as E .

Each entity $e \in E$ is linked to a Wikidata QID, from which we obtain relevant metadata (e.g., country of origin, entity type). The following sections

define each metric in detail.

4.1 Granularity

This metric captures the model’s ability to reference culturally specific concepts. We define **granularity** as the level of cultural detail reflected in the predicted entities. Inspired by a prior work on semantic abstraction in human descriptions (Berger and Ponti, 2025), we assign a granularity score to each entity based on a predefined cultural scale (e.g. music genre vs. song title, see Appendix B.2), with higher values indicating finer granularity. Let $\text{Gran}(e)$ be the granularity level of entity e . We compute:

$$\text{Granularity} = \frac{1}{|E|} \sum_{e \in E} \text{Gran}(e), \quad (1)$$

4.2 Diversity

Diversity measures the breadth of cultural coverage by counting the number of unique cultural entities generated and is inspired by and adapted from work by Berger and Ponti (2025). As shown in Figure 2, the prompt asks about books, and the metric calculates the number of distinct book-related entities mentioned. Although the English response contains three extracted entities, one is repeated, so the resulting diversity score is 2. Higher diversity indicates that the model is capable of expressing a wider range of cultural knowledge for a given language and topic.

4.3 Culture Specificity

Culture specificity evaluates whether the entities in the generated text are aligned with the cultural context specified in the prompt. Figure 2 illustrates a case where the prompt inquires about the US. The culture specificity score is calculated by computing the proportion of entities associated with the United States among all extracted entities. This metric reflects the model’s ability to generate culturally appropriate and contextually aligned content. Its design is motivated by earlier studies on cultural awareness and bias (Wang et al., 2024; Liu et al., 2024b).

4.4 Culture Consensus

This metric is motivated by the observations of Hedderich et al. (2025), who showed that LLMs prompted with the same query in different languages often yield inconsistent outputs, highlighting the need to evaluate whether models preserve

shared cultural understanding across languages. **Culture consensus** quantifies the agreement between model outputs for the same topic and cultural context setting across different languages. Let Q_A and Q_B denote the sets of Wikidata QIDs extracted from responses to the same prompt in languages A and B , respectively. We compute their overlap using Jaccard similarity:

$$\text{Consensus}(A, B) = \frac{|Q_A \cap Q_B|}{|Q_A \cup Q_B|}. \quad (2)$$

Higher consensus implies more stable cultural grounding across languages, suggesting that the model maintains consistent cultural knowledge despite surface linguistic variation.

5 Experimental Setups

5.1 Models

We evaluate our framework on seven multilingual language models, chosen to represent a diverse range of capabilities, resource requirements, and geographic origins. Our selection includes three open-source models developed across different regions: **Llama-3.1-8B-Instruct** (US) (Grattafiori et al., 2024), **Mistral-7B-Instruct-v0.1** (France) (Jiang et al., 2023), and **Qwen2.5-7B-Instruct** (China) (Yang et al., 2024). To complement these, we also include two proprietary large-scale models: **ChatGPT-4o-mini** (US) and **DeepSeek-V3** (China) (Liu et al., 2024a), which are widely used but not openly available or easily deployable on local infrastructure. In addition, we evaluate **Llama-3.3-70B-Instruct** (Grattafiori et al., 2024) and **aya-expansion-8b** (Canada) (Üstün et al., 2024). The inclusion of both Llama3 variants enables a controlled comparison of model size, while Aya allows us to assess cross-lingual generalization in a model explicitly trained for multilinguality.

5.2 Languages

We aim at a balanced and diverse coverage of major cultural regions and continents. Therefore, we select 13 languages as the primary focus of our study: Arabic, English, German, Hindi, Italian, Japanese, Korean, Persian, Simplified Chinese, Spanish, Thai, Traditional Chinese, and Turkish. The chosen languages span multiple language families (e.g., Indo-European, Sino-Tibetan), writing systems (e.g., Kanji, Devanagari), and geographic regions (for example, Middle East). Table 2 provides an overview of the selected languages and

Type	Prompt
Neutral	Write a short text (around 50 words) about a person having dinner.
Explicit	Write a short text (around 50 words) about a person from {country} having dinner.
Implicit	Write a short text (around 50 words) about {name} having dinner.

Table 1: English prompt templates for the *food* category under three conditions: **Neutral**, which makes no cultural reference; **Explicit** and **Implicit** include a hint for the referred cultural context. Full prompt templates for all categories are provided in Appendix B.1.

Lang	Countries/Regions
ar	United Arab Emirates
en	United States, United Kingdom, Canada, Australia, Nigeria
de	Germany
es	Mexico, Spain, Argentina
fa	Iran
hi	India
it	Italy
ja	Japan
ko	South Korea
th	Thailand
tr	Turkey
zh	China
zh-tw	Taiwan

Table 2: Prompt languages and country/region mentioned countries/regions in our experiments. Country/region names are color-coded by geographical region: East Asia, South Asia, Middle East, Europe, North America, Latin America, Australia and Africa.

their corresponding regions²

5.3 Prompts

We design three categories of prompts to elicit culturally grounded responses from language models: **explicit contextualized prompts**, where the cultural background is explicitly mentioned (e.g. “about a person from Spain”), **implicit contextualized prompts**, where a culturally associated name is mentioned (e.g. “about Juan”, where “Juan” is a popular name in Spain) and the **neutral prompts** with no mention of nationality, ethnicity, or location, allowing us to assess cultural preferences expressed by the model without direct cues. Table 1 provides examples of contextualized and neutral prompt templates.

All prompts are constructed across six topics:

²In addition, we initially explored two low-resource African languages: Hausa and Yoruba. However, during experimentation, we found that open-source models produced text of very poor quality in these two languages, often entirely unrelated to the prompts (see Appendix B). As a result, we set aside experiments involving native African languages as future work.

food, beverages, clothing, books, music, and transportation. These categories were chosen because they reflect culturally meaningful yet concrete aspects of daily life that are comparable across languages.

To ensure linguistic accuracy, all prompt templates are reviewed by native speakers to validate both meaning and fluency. For Spanish, the prompts were checked by native speakers from both Latin America and Europe to account for regional variation. For morphologically-rich languages like Turkish, we design templates that keep the placeholder position grammatically neutral to allow automatic template filling.

During pilot experiments, we observed notable inconsistencies in language alignment for certain models. Specifically, Mistral frequently responded in English even when prompted in other languages, while Qwen exhibited a strong tendency to reply in Chinese, reaching up to 80% of responses. Such behavior was not observed in other models under the same conditions. To mitigate this, we appended an additional instruction to the prompts used with these models to explicitly enforce language consistency. Details of prompts are provided in Appendix B.

6 Results and Analysis

Our evaluation is based on a total of 1,716 unique prompts, covering 13 languages, 6 cultural topics, 19 countries, and both contextualized and neutral settings. Each prompt is used to generate 500 responses per model, resulting in over 85.8 million generated texts. MAKIEVAL identifies a total of 24,042 matched cultural entities (details in Appendix C.3), which form the basis for computing our four proposed metrics.

In the remainder of this work, we explore the factors that influence cultural granularity, diversity, specificity and consensus, highlighting the complex interplay between model, language as well as cultural context.




Figure 3: Average granularity ratio for Chinese, English, and Turkish prompts across all models. Each point represents the average granularity score aggregated over six cultural topics. Full results are provided in Appendix D.




Figure 4: Diversity metric comparing prompts with culture context India in the native language Hindi vs. other languages for Llama3. Each row shows a separate cultural topic.

6.1 Granularity

Our analysis reveals that, for each model, the average granularity is mostly the same given the same prompt language, independent of the referenced country/cultural context. We show the results for Chinese, English and Turkish as typical representatives in Figure 3, with all results given in Appendix D. Individual outliers exist (e.g., Mistral prompted in Turkish with Germany as context), but no context-specific pattern is visible. This suggests that each model maintains a language-specific baseline for cultural granularity that remains stable across cultural contexts.

Differences exist when moving with the same model to a different prompt language. E.g., English and Turkish have lower granularity levels for most models compared to Chinese. The observed patterns may stem from inherent linguistic differences in lexical specificity when describing cultural concepts, where Chinese, e.g., might employ more fine-grained expressions. Or alternatively, these differences might reflect limitations in the models’ multilingual capabilities.

6.2 Diversity

The diversity metric varies widely depending on the setting, ranging from a handful of distinct entities per topic to over a hundred. The differences depend on both the prompt language and the model, with Llama3 having the highest mean diversity of 36.65 and DeepSeek the lowest with 14.62. For DeepSeek, when prompted in English on the book topic for the United States, we observe a diversity score close to 1, with the model producing the entity *To Kill a Mockingbird* in 499 out of 500 generations.

Online texts in the native language of a country, used to train models, likely contain a diverse and maybe even less stereotypical representation of a country’s culture compared to texts in other languages. We, therefore, explore whether language models exhibit stronger cultural diversity when prompted in the native language of the country being described. On average, the diversity is higher for a country when prompted in the country’s native language (26.518 vs. 23.617). Figure 4 shows India and Hindi as examples. There is, however, no guarantee that using the native language increases diversity, with Appendix E also showing many counterexamples across models, languages and topics. For English and Spanish, languages spoken natively in different countries, the diversity also differs by country. This shows that cultural diversity representation goes beyond simply the prompt language.

6.3 Cultural Specificity

Taking Llama3 as an example, shown in Figure 5 left, we can see that the neutral prompt (without mention of a cultural context or country) has a high cultural specificity when matching countries with their native language (e.g., Chinese-China and

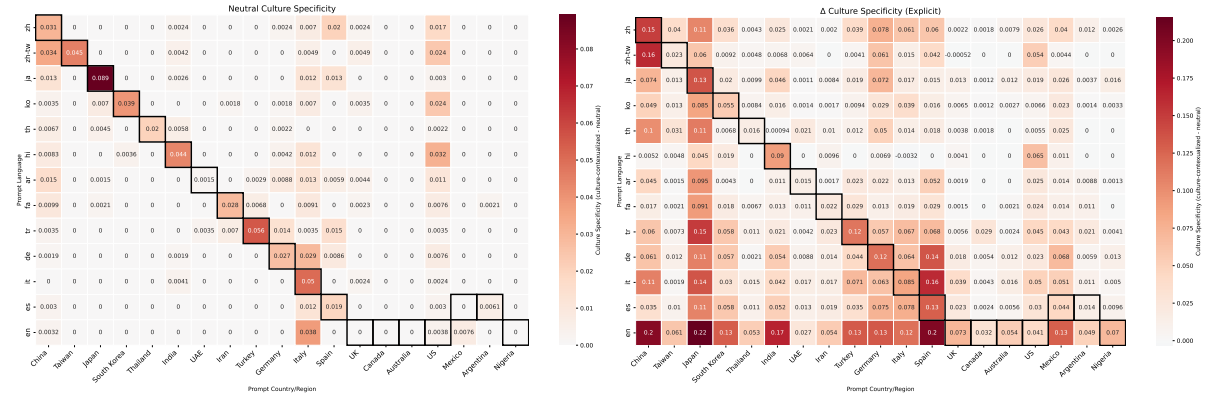


Figure 5: Culture specificity across country-language pairs for Llama3 without cultural context (neutral prompt, left) and with explicit context (right). Black boxes mark matching country and native language.

Model	Culture Consensus
Aya	0.1839 ± 0.0030
ChatGPT	0.2036 ± 0.0015
DeepSeek	0.2013 ± 0.0015
Llama3	0.1813 ± 0.0010
Llama3-70B	0.2097 ± 0.0011
Mistral	0.1106 ± 0.0041
Qwen	0.1623 ± 0.0018

Table 3: Culture consensus scores for different models averaged over all topics and languages (mean ± variance).

Turkish-Turkey). This suggests that language inherently encodes cultural signals. The model’s outputs are, however, not limited to this language-country connection with entities from other countries also appearing in the outputs, especially from China, Italy and the United States.

Prompting with the explicit cultural context increases the cultural specificity in most cases (Figure 5 right). Implicit prompts (that do not mention the country explicitly) in LLaMA3 produce weaker gains than explicit ones (Figure 21 in Appendix), indicating reliance on overt cultural cues. Crucially, English prompts consistently achieve the highest specificity gains when adding cultural context across nearly all countries and models (see Appendix F). This suggests that LLMs are more sensitive to cultural signals in English, likely due to the disproportionate amount of English data in their pretraining corpora. This further underscores the importance of multilingual evaluation, as cultural awareness is not only model-dependent but also deeply shaped by language.

6.4 Cultural Consensus

Cultural consensus measures for a model how similar the cultural entities are across different prompt languages. Note that this is computed for all cultural contexts, not only that of the native language. A high consensus might thus be an indicator that the model has a representation of the cultures that is independent of the specific prompt language.

As seen in Table 3, model size seems to affect consensus, with the group of three larger models having a similar consensus value of around 0.2. Model size is, however, not the only factor, with the difference in consensus being high among the smaller models of similar size. The figures in Appendix G, provide a more detailed picture with large differences of consensus across language pairs. Taking ChatGPT on the food topic as an example in Figure 6, we see that there is a high consensus between pairs of European languages as well as between Japanese and Korean. We hypothesize that this intra-regional alignment could be caused by shared perspectives on their own countries and those abroad, which then manifests in large-scale web or multilingual corpora through food blogs, reviews, or regional discussions. This effect suggests that cultural awareness in language models may be influenced not only by language itself but also by the latent cultural connections among the communities represented by those languages.

7 Discussion

The design of the metrics in MAKIEVAL is informed by insights from psycholinguistics and linguistics, as well as prior work in NLP, and our experiments demonstrate systematic differences

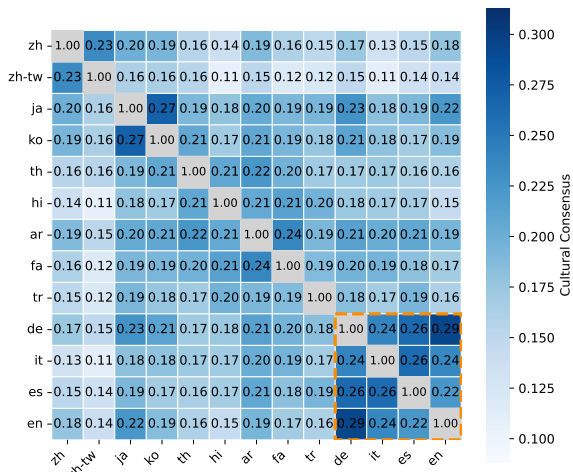


Figure 6: Cultural consensus between language pairs (over all countries) for the food topic for ChatGPT.

across languages and models. However, how these differences should be interpreted and applied in practice remains an open question that merits further discussion. MAKIEVAL is not intended as a prescriptive benchmark in which higher values are inherently preferable across all dimensions. Instead, it provides a set of complementary measures to capture different aspects of how LLMs represent and express cultural knowledge.

The relevance and desirability of these measures vary with the characteristics of the downstream task and the interpretation of the proposed metrics requires careful consideration. For example, high specificity and granularity may be beneficial in settings where interlocutors share a cultural background, while lower values may be more effective in intercultural communication where common ground cannot be assumed. Similarly, cross-lingual consensus may indicate robustness in knowledge-intensive applications such as multilingual question answering and cross-lingual understanding (Chen et al., 2022b; Qi et al., 2023; Wang et al., 2025), whereas divergence can be advantageous in tasks that benefit from creativity or cultural variation, including storytelling or adaptation to local contexts. Given the diversity of generation tasks, no single value or direction of change should be considered universally optimal. Instead, the task-specific needs should be considered along the the proposed cultural dimensions.

8 Conclusion

We present MAKIEVAL, an automatic, scalable and multilingual framework for evaluating cul-

tural awareness in LLMs. By combining prompt-controlled generation, structured entity extraction, Wikidata grounding, and four culture metrics, our framework enables analysis of how models encode and express cultural knowledge across different languages, countries/regions, and topics.

Our results highlight that cultural awareness in LLMs is more complex than simple descriptions like a general bias to Western culture. Rather, it manifests in different forms, like the granularity, specificity and diversity of cultural entities. Besides general model capabilities, we show that the prompting language plays a crucial role in evoking these cultural differences, even when no explicit cultural context is given. The sensitivity to cultural signals in English, the connection between countries and their local languages, and the cultural consensus that links regional areas emphasize the need for multilingual and multicontextual evaluation of cultural awareness.

To foster further analysis of the cultural awareness of LLMs, we will release the full dataset of generated outputs and cultural entities for all model, language, country and topic combinations. Our work also provides guidance for future research in identifying how these factors have been picked up by the models during training and alignment, and how culture is represented internally in the models. More broadly, our framework offers a generalizable paradigm for multilingual evaluation, our findings reveal systematic disparities that call for rethinking prompting and training beyond English, and our metrics provide actionable signals for improving fairness and inclusivity in global LLM applications. We hope this work contributes to more equitable evaluation and deployment of language technologies worldwide.

Limitations

While MAKIEVAL enables scalable multilingual evaluation of cultural awareness, it has several limitations.

First, our experiments are limited to a subset of high- and mid-resource languages. Although Wikidata provides labels in many low-resource languages, we excluded low-resource languages due to the low generation quality from current language models in these settings.

Second, our method relies on metadata retrieved from Wikidata to match and contextualize cultural entities. As an open, community-curated knowl-

edge base, Wikidata may contain incomplete or inaccurate information. In practice, we observe that such cases are rare and have limited impact on our overall evaluation, though they may introduce noise in entity-level metrics such as specificity or consensus.

Finally, MAKIEVAL primarily targets concrete cultural expression—entities such as food, books, and music that are grounded and observable. It is not designed to evaluate abstract cultural constructs, such as moral reasoning, political ideology, or values, which remain important directions for future work.

Ethical Concern We do not foresee any ethical concerns associated with this work. All analyses were conducted using publicly available models. No private or sensitive information was used. The texts we studied were model-generated and do not necessarily represent our views. Additionally, we will release our code, produced data and documentations to support transparency and reproducibility.

Acknowledgments

We thank the members of MaiNLP for their valuable feedback on this project, especially Verena Blaschke, Yupei Du, Florian Echin, Rob van der Goot, Bolei Ma, Elena Senger, and Soh-Eun Shim. We are further grateful to a number of these colleagues, as well as to Fatima Abdulla Mohammed Obaid Almatrooshi, Tianshi Feng, Paravee Jungbauer, Neysa Kaustubh, Abdullatif Köksal, Yuchen Li, Ali Modarressi, and Jing’an Qian for their assistance in verifying the language-specific texts studied in this paper.

This research was supported by the ERC Consolidator Grant DIALECT 101043235. We gratefully acknowledge that experiments involving API calls to GPT-4o-mini were supported by a compute grant from OpenAI.

References

Utkarsh Agarwal, Kumar Tanmay, Aditi Khandelwal, and Monojit Choudhury. 2024. [Ethical reasoning and moral value alignment of LLMs depend on the language we prompt them in](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 6330–6340, Torino, Italia. ELRA and ICCL.

Yejin Bang, Delong Chen, Nayeon Lee, and Pascale Fung. 2024. [Measuring political bias in large lan-](#)

[guage models: What is said and how it is said](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11142–11159, Bangkok, Thailand. Association for Computational Linguistics.

Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.

Uri Berger and Edoardo Ponti. 2025. [Cross-lingual and cross-cultural variation in image descriptions](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 9453–9465, Albuquerque, New Mexico. Association for Computational Linguistics.

Shaily Bhatt and Fernando Diaz. 2024. [Extrinsic evaluation of cultural competence in large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 16055–16074, Miami, Florida, USA. Association for Computational Linguistics.

Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is power: A critical survey of “bias” in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.

Wallace Chafe. 1994. *Discourse, consciousness, and time: The flow and displacement of conscious experience in speaking and writing*. University of Chicago Press.

Beiduo Chen, Wu Guo, Bin Gu, Quan Liu, and Yongchao Wang. 2022a. [Multi-level contrastive learning for cross-lingual alignment](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2022, Virtual and Singapore, 23-27 May 2022*, pages 7947–7951. IEEE.

Beiduo Chen, Wu Guo, Quan Liu, and Kun Tao. 2022b. [Feature aggregation in zero-shot cross-lingual transfer using multilingual BERT](#). In *26th International Conference on Pattern Recognition, ICPR 2022, Montreal, QC, Canada, August 21-25, 2022*, pages 1428–1435. IEEE.

Beiduo Chen, Jun-Yu Ma, Jiajun Qi, Wu Guo, Zhen-Hua Ling, and Quan Liu. 2022c. [USTC-NELSLIP at SemEval-2022 task 11: Gazetteer-adapted integration network for multilingual complex named entity recognition](#). In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1613–1622, Seattle, United States. Association for Computational Linguistics.

- Myra Cheng, Esin Durmus, and Dan Jurafsky. 2023. [Marked personas: Using natural language prompts to measure stereotypes in language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1504–1532, Toronto, Canada. Association for Computational Linguistics.
- Yu Ying Chiu, Liwei Jiang, Bill Yuchen Lin, Chan Young Park, Shuyue Stella Li, Sahithya Ravi, Mehar Bhatia, Maria Antoniak, Yulia Tsvetkov, Vered Shwartz, and Yejin Choi. 2024. [Cultural-bench: a robust, diverse and challenging benchmark on measuring the \(lack of\) cultural knowledge of llms](#). *Preprint*, arXiv:2410.02677.
- Ameet Deshpande, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, and Karthik Narasimhan. 2023. [Toxicity in chatgpt: Analyzing persona-assigned language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1236–1270, Singapore. Association for Computational Linguistics.
- Jean-Marc Dewaele and Aneta Pavlenko. 2003. Productivity and lexical diversity in native and non-native speech: A study of cross-cultural effects. *Effects of the second language on the first*, 3:120.
- Sara Ghaboura, Ketan Pravin More, Ritesh Thawkar, Wafa Al Ghallabi, Omkar Thawakar, Fahad Shahbaz Khan, Hisham Cholakkal, Salman Khan, and Rao Muhammad Anwer. 2025. [Time travel: A comprehensive benchmark to evaluate LMMs on historical and cultural artifacts](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 23627–23641, Vienna, Austria. Association for Computational Linguistics.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Michael A. Hedderich, Anyi Wang, Raoyuan Zhao, Florian Eichin, Jonas Fischer, and Barbara Plank. 2025. [What’s the difference? supporting users in identifying the effects of prompt and model changes through token patterns](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 20093–20123, Vienna, Austria. Association for Computational Linguistics.
- Tianyi Hu, Maria Maistro, and Daniel Hershcovich. 2024. [Bridging cultures in the kitchen: A framework and benchmark for cross-cultural recipe retrieval](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1068–1080, Miami, Florida, USA. Association for Computational Linguistics.
- Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denuyl. 2020. [Social biases in NLP models as barriers for persons with disabilities](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5491–5501, Online. Association for Computational Linguistics.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, and 1 others. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Ming Jiang and Mansi Joshi. 2024. [CPopQA: Ranking cultural concept popularity by LLMs](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 615–630, Mexico City, Mexico. Association for Computational Linguistics.
- Jiho Jin, Jiseon Kim, Nayeon Lee, Haneul Yoo, Alice Oh, and Hwaran Lee. 2024. [KoBBQ: Korean bias benchmark for question answering](#). *Transactions of the Association for Computational Linguistics*, 12:507–524.
- Karen Sparck Jones and Julia R Galliers. 1995. Evaluating natural language processing systems: An analysis and review.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, and 1 others. 2024a. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Chen Cecilia Liu, Iryna Gurevych, and Anna Korhonen. 2024b. Culturally aware and adapted nlp: A taxonomy and a survey of the state of the art. *arXiv preprint arXiv:2406.03930*.
- Bolei Ma, Xinpeng Wang, Tiancheng Hu, Anna-Carolina Haensch, Michael A. Hedderich, Barbara Plank, and Frauke Kreuter. 2024. [The potential and challenges of evaluating attitudes, opinions, and values in large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 8783–8805, Miami, Florida, USA. Association for Computational Linguistics.
- Junho Myung, Nayeon Lee, Yi Zhou, Jiho Jin, Rifki Putri, Dimosthenis Antypas, Hsuvas Borkakoty, Eunsu Kim, Carla Perez-Almendros, Abinew Ali Ayele, and 1 others. 2024. Blend: A benchmark for llms on everyday knowledge in diverse cultures and languages. *Advances in Neural Information Processing Systems*, 37:78104–78146.
- Tarek Naous, Michael J Ryan, Alan Ritter, and Wei Xu. 2024. [Having beer after prayer? measuring cultural bias in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16366–16393, Bangkok, Thailand. Association for Computational Linguistics.

- Shravan Nayak, Kanishk Jain, Rabiul Awal, Siva Reddy, Sjoerd Van Steenkiste, Lisa Anne Hendricks, Karolina Stanczak, and Aishwarya Agrawal. 2024. [Benchmarking vision language models for cultural understanding](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5769–5790, Miami, Florida, USA. Association for Computational Linguistics.
- Malvina Nikandrou, Georgios Pantazopoulos, Nikolas Vitsakis, Ioannis Konstas, and Alessandro Suglia. 2025. [CROPE: Evaluating in-context adaptation of vision and language models to culture-specific concepts](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7917–7936, Albuquerque, New Mexico. Association for Computational Linguistics.
- Shota Onohara, Atsuyuki Miyai, Yuki Imajuku, Kazuki Egashira, Jeonghun Baek, Xiang Yue, Graham Neubig, and Kiyoharu Aizawa. 2025. [JMMMU: A Japanese massive multi-discipline multimodal understanding benchmark for culture-aware evaluation](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 932–950, Albuquerque, New Mexico. Association for Computational Linguistics.
- OpenAI. 2025. [Openai o3 and o4-mini](https://openai.com/index/introducing-o3-and-o4-mini/). <https://openai.com/index/introducing-o3-and-o4-mini/>. Accessed: 2025-05-10.
- Siddhesh Pawar, Junyeong Park, Jiho Jin, Arnav Arora, Junho Myung, Srishti Yadav, Faiz Ghifari Haznitrana, Inhwa Song, Alice Oh, and Isabelle Augenstein. 2024. [Survey of cultural awareness in language models: Text and beyond](#). *Preprint*, arXiv:2411.00860.
- Siddhesh Pawar, Junyeong Park, Jiho Jin, Arnav Arora, Junho Myung, Srishti Yadav, Faiz Ghifari Haznitrana, Inhwa Song, Alice Oh, and Isabelle Augenstein. 2025. [Survey of cultural awareness in language models: Text and beyond](#). *Computational Linguistics*, pages 1–96.
- Ryan E Peters, Theres Grüter, and Arielle Borovsky. 2018. Vocabulary size and native speaker self-identification influence flexibility in linguistic prediction among adult bilinguals. *Applied Psycholinguistics*, 39(6):1439–1469.
- Jirui Qi, Raquel Fernández, and Arianna Bisazza. 2023. [Cross-lingual consistency of factual knowledge in multilingual language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10650–10666, Singapore. Association for Computational Linguistics.
- Florian Schneider and Sunayana Sitaram. 2024. [M5 – a diverse benchmark to assess the performance of large multimodal models across multilingual and multi-cultural vision-language tasks](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4309–4345, Miami, Florida, USA. Association for Computational Linguistics.
- Andrea Seveso, Daniele Poterì, Edoardo Federici, Mario Mezzananza, and Fabio Mercorio. 2025. [ITALIC: An Italian culture-aware natural language benchmark](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1469–1478, Albuquerque, New Mexico. Association for Computational Linguistics.
- Weiyan Shi, Ryan Li, Yutong Zhang, Caleb Ziems, Sunny Yu, Raya Horesh, Rogério Abreu De Paula, and Diyi Yang. 2024. [CultureBank: An online community-driven knowledge base towards culturally aware language technologies](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4996–5025, Miami, Florida, USA. Association for Computational Linguistics.
- Shivalika Singh, Angelika Romanou, Clémentine Fourrier, David Ifeoluwa Adelani, Jian Gang Ngui, Daniel Vila-Suero, Peerat Limkonchotiwat, Kelly Marchisio, Wei Qi Leong, Yosephine Susanto, Raymond Ng, Shayne Longpre, Sebastian Ruder, Wei-Yin Ko, Antoine Bosselut, Alice Oh, Andre Martins, Leshem Choshen, Daphne Ippolito, and 4 others. 2025. [Global MMLU: Understanding and addressing cultural and linguistic biases in multilingual evaluation](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 18761–18799, Vienna, Austria. Association for Computational Linguistics.
- Ahmet Üstün, Viraat Aryabumi, Zheng Yong, Wei-Yin Ko, Daniel D’souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, Freddie Vargus, Phil Blunsom, Shayne Longpre, Niklas Muennighoff, Marzieh Fadaee, Julia Kreutzer, and Sara Hooker. 2024. [Aya model: An instruction fine-tuned open-access multilingual language model](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15894–15939, Bangkok, Thailand. Association for Computational Linguistics.
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.
- Yixin Wan, George Pu, Jiao Sun, Aparna Garimella, Kai-Wei Chang, and Nanyun Peng. 2023. [“kelly is a warm person, joseph is a role model”](#): Gender biases in LLM-generated reference letters. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3730–3748, Singapore. Association for Computational Linguistics.
- Mingyang Wang, Heike Adel, Lukas Lange, Yihong Liu, Ercong Nie, Jannik Strötgen, and Hinrich Schuetze.

2025. [Lost in multilinguality: Dissecting cross-lingual factual inconsistency in transformer language models](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5075–5094, Vienna, Austria. Association for Computational Linguistics.

Wenxuan Wang, Wenxiang Jiao, Jingyuan Huang, Ruyi Dai, Jen-tse Huang, Zhaopeng Tu, and Michael Lyu. 2024. [Not all countries celebrate thanksgiving: On the cultural dominance in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6349–6384, Bangkok, Thailand. Association for Computational Linguistics.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, and 1 others. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.

Binwei Yao, Ming Jiang, Tara Bobinac, Diyi Yang, and Junjie Hu. 2024. [Benchmarking machine translation with cultural awareness](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13078–13096, Miami, Florida, USA. Association for Computational Linguistics.

Arda Yüksel, Abdullatif Köksal, Lütfi Kerem Senel, Anna Korhonen, and Hinrich Schuetze. 2024. [TurkishMMLU: Measuring massive multitask language understanding in Turkish](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7035–7055, Miami, Florida, USA. Association for Computational Linguistics.

Li Zhou, Taelin Karidi, Wanlong Liu, Nicolas Garneau, Yong Cao, Wenyu Chen, Haizhou Li, and Daniel Hershcovich. 2025. [Does mapo tofu contain coffee? probing LLMs for food-related cultural knowledge](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 9840–9867, Albuquerque, New Mexico. Association for Computational Linguistics.

A Computational Resources and AI Assistance

We checked the licenses of all the models and data used, which are publicly available resources.

Computational Resources We generate the texts with temperature = 0.7, top- p = 0.9, and top- k = 10 for all LLMs that run locally. All models were executed on NVIDIA A100 GPUs, except for the LLaMA3-70B model, which was run on an H200 GPU. The DeepSeek and ChatGPT models were accessed via their respective public APIs.

AI Assistance The authors acknowledge the use of ChatGPT solely for correcting grammatical errors, enhancing the coherence of the final manuscript, and providing assistance with coding.

B Prompt Design

B.1 Prompt Template for Text Generation

We use structured templates to generate prompts across six topics and 13 languages. To ensure linguistic accuracy and cultural appropriateness, we invited native speakers to review and validate the prompts. A total of 15 native speakers participated in this validation process, with at least one annotator per language. Their language backgrounds include German, Spanish (Spain), Spanish (Mexico), English (American), English (British), Turkish, Persian, Arabic, Hindi, Japanese, Korean, Traditional Chinese (Taiwan), Thai, Chinese(Mainland) and Italian.

To avoid redundancy, we present prompt templates for the food topic in Table 7 here only. Templates for other topics share the same structural format and are available in the released code.

During our experiments, we observed a notable mismatch between the prompt language and the generated response language in the Mistral and Qwen models. To ensure accurate evaluation, we appended a language-specific instruction as listed in Table 4 to the end of the original prompt when querying these models. This addition effectively resolved the misalignment issue, enabling the models to generate responses in the intended language.

Language	Instruction Prompt
ar	يرجى الرد باللغة العربية.
de	Bitte antworten Sie auf Deutsch.
en	Please answer in English.
es	Por favor, responde en español.
fa	لطفاً به زبان فارسی پاسخ دهید.
hi	कृपया हिंदी में उत्तर दें।
it	Si prega di rispondere in italiano.
ja	日本語で教えてください。
ko	한국어로 대답해주세요.
th	กรุณาตอบเป็นภาษาไทย
tr	Lütfen Türkçe cevap verin.
zh	请用中文作答。
zh-tw	請用繁體中文作答。

Table 4: Language-specific answer instructions.

Lang.	Prompt (Neutral, Explicit, Implicit-Male, Implicit-Female)
ar	اكتب نصاً قصيراً (حوالي 05 كلمة) عن شخص يتناول العشاء. يتناول العشاء. {country} اكتب نصاً قصيراً (حوالي 05 كلمة) عن شخص من يتناول العشاء. {name_male} اكتب نصاً قصيراً (حوالي 05 كلمة) عن تتناول العشاء. {name_female} اكتب نصاً قصيراً (حوالي 05 كلمة) عن
de	Schreiben Sie einen kurzen Text (ca. 50 Wörter) über eine Person, die Abendessen isst. Schreiben Sie einen kurzen Text (ca. 50 Wörter) über eine Person aus {country}, die Abendessen isst. Schreiben Sie einen kurzen Text (ca. 50 Wörter) über {name_male}, der Abendessen isst. Schreiben Sie einen kurzen Text (ca. 50 Wörter) über {name_female}, die Abendessen isst.
en	Write a short text (around 50 words) about a person having dinner. Write a short text (around 50 words) about a person from {country} having dinner. Write a short text (around 50 words) about {name_male} having dinner. Write a short text (around 50 words) about {name_female} having dinner.
es	Escribe un texto corto (de unas 50 palabras) sobre una persona que está cenando. Escribe un texto corto (de unas 50 palabras) sobre una persona de {country} que está cenando. Escribe un texto corto (de unas 50 palabras) sobre {name_male} que está cenando. Escribe un texto corto (de unas 50 palabras) sobre {name_female} que está cenando.
fa	یک متن کوتاه (حدود ۰۵ کلمه) درباره فردی که در حال صرف شام است بنویس. که در حال صرف شام است بنویس. {country} یک متن کوتاه (حدود ۰۵ کلمه) درباره فردی از کشور که در حال صرف شام است بنویس {name_male} یک متن کوتاه (حدود ۰۵ کلمه) درباره که در حال صرف شام است بنویس. {name_female} یک متن کوتاه (حدود ۰۵ کلمه) درباره
hi	एक व्यक्ति जो रात का खाना खा रहा है, उसके बारे में लगभग 50 शब्दों का एक संक्षिप्त अनुच्छेद लिखें। एक व्यक्ति जो {country} से है और रात का खाना खा रहा है, उसके बारे में लगभग 50 शब्दों का एक संक्षिप्त अनुच्छेद लिखें। {name_male} जो रात का खाना खा रहा है, उसके बारे में लगभग 50 शब्दों का एक संक्षिप्त अनुच्छेद लिखें। {name_female} जो रात का खाना खा रही है, उसके बारे में लगभग 50 शब्दों का एक संक्षिप्त अनुच्छेद लिखें।
it	Scrivi un breve testo (circa 50 parole) su una persona che sta cenando. Scrivi un breve testo (circa 50 parole) su una persona proveniente da {country} che sta cenando. Scrivi un breve testo (circa 50 parole) su {name_male} che sta cenando. Scrivi un breve testo (circa 50 parole) su {name_female} che sta cenando.
ja	夕食を食べている人について、50語程度の短い文章を書いてください。 {country}出身の人が夕食を食べている様子について、50語程度の短い文章を書いてください。 {name_male}が夕食を食べている様子について、50語程度の短い文章を書いてください。 {name_female}が夕食を食べている様子について、50語程度の短い文章を書いてください。
ko	어떤 사람이 저녁 식사를 하고 있는 모습을 약 50단어 분량의 짧은 글로 작성하세요. {country} 출신의 어떤 사람이 저녁 식사를 하고 있는 모습을 약 50단어 분량의 짧은 글로 작성하세요. {name_male}이 저녁 식사를 하고 있는 모습을 약 50단어 분량의 짧은 글로 작성하세요. {name_female}이 저녁 식사를 하고 있는 모습을 약 50단어 분량의 짧은 글로 작성하세요.
th	เขียนข้อความสั้นประมาณ 50 คำเกี่ยวกับบุคคลที่กำลังกินอาหารเย็นอยู่ เขียนข้อความสั้นประมาณ 50 คำเกี่ยวกับบุคคลจากประเทศ {country} ที่กำลังกินอาหารเย็นอยู่ เขียนข้อความสั้นประมาณ 50 คำเกี่ยวกับ {name_male} ที่กำลังกินอาหารเย็นอยู่ เขียนข้อความสั้นประมาณ 50 คำเกี่ยวกับ {name_female} ที่กำลังกินอาหารเย็นอยู่
tr	Akşam yemeği yiyen bir kişi hakkında (yaklaşık 50 kelimelik) kısa bir metin yazın. {country} vatandaşı bir kişinin akşam yemeği yemesi hakkında (yaklaşık 50 kelimelik) kısa bir metin yazın. Akşam yemeği yiyen {name_male} isimli bir kişi hakkında (yaklaşık 50 kelimelik) kısa bir metin yazın. Akşam yemeği yiyen {name_female} isimli bir kişi hakkında (yaklaşık 50 kelimelik) kısa bir metin yazın.
zh	写一段简短的文字 (大约50个字), 描述一个人在吃晚饭。 写一段简短的文字 (大约50个字), 描述一个来自{country}的人在吃晚饭。 写一段简短的文字 (大约50个字), 描述{name_male}在吃晚饭。 写一段简短的文字 (大约50个字), 描述{name_female}在吃晚饭。
zh-tw	寫一段簡短的文字 (約 50 個字), 描述某人正在吃晚餐。 寫一段簡短的文字 (約 50 個字), 描述來自 {country} 的某人正在吃晚餐。 寫一段簡短的文字 (約 50 個字), 描述{name_male} 正在吃晚餐。 寫一段簡短的文字 (約 50 個字), 描述{name_female} 正在吃晚餐。

Figure 7: Neutral, explicit, and implicit prompts (food topic) across 13 languages.

B.2 Prompt for Entity Extraction

All prompts for entity extraction follow this format:

```
Help me extract the words or phrases in
the given text which are {topic-specific
types}.
Answer in a dictionary format with the
type of the extracted text labeled.
Do not extract redundant entities or
provide explanations.
Only output the dictionary in this
format:
{"<extracted_text>": "type",
...}
Text: {response}
```

And based on the topic, as shown in Table 5, the system dynamically switches to different extraction branches. In addition to extracting topic-specific entities, we also include named entities such as person names and geographic locations to support more comprehensive downstream analyses.

Furthermore, we require that each extracted entity be assigned a granularity label. While we initially considered leveraging Wikidata to define these granular levels, the hierarchy was often inconsistent and difficult to generalize across languages and cultures. As a result, we employed ChatGPT-4o-mini to infer the granularity of each entity. For example, food-related entities were categorized into labels such as dish, dish category, specific ingredient, and ingredient category. To quantify overall granularity, we assign a score of 1 to entities with specific references (e.g., spaghetti, coriander) and a score of 0 to more general categories (e.g., pasta dishes, herbs). The final granularity score for a given set of entities is computed as the average of these binary scores. More details can be found in the prompt design in Table 5.

B.3 Generation Quality in Low-Resource African Languages

To expand our evaluation beyond high-resource languages, we initially included Hausa and Yoruba in early-stage experiments. However, we observed that open-source models consistently produced low-quality outputs in these two languages. As shown in Table 6, many generations were either semantically irrelevant, grammatically incoherent, or entirely disconnected from the prompt intent. Due to this, we excluded Hausa and Yoruba from the main evaluation.

C Framework Evaluation

C.1 Performance of Entity Extraction

We randomly sampled 50 sentences for each of six topics across four languages: English, Chinese, Arabic, and Hindi, to manually evaluate the performance of ChatGPT-4o-mini on entity extraction. Each language was annotated by one native speaker of that language. All annotators were native speakers with a graduate-level education background. While only one annotator was involved per language, we ensured careful guidelines and consistency checks to support evaluation reliability. As shown in Table 8 below, the F1 scores across all languages are consistently high, demonstrating the reliability of the entity extraction process. Additionally, as can be seen in Table 9, extraction accuracy is evaluated by category for English with 50 samples from each domain. Although performance varies across different domains, the overall results remain consistently strong.

Our annotators noted that missed entities were rare, but in some cases, the model incorrectly identified non-target items as entities. A common error observed across all four languages in the food topic was that utensils were sometimes mistakenly labeled as food entities. While this type of misidentification, such as mistaking utensils for food was observed across languages, which can actually be fixed in our pipeline. Specifically, our framework’s Wikidata-based disambiguation step is capable of filtering out such incorrect entities, thereby mitigating their impact during downstream evaluation.

C.2 Performance of Wikidata Search

Wikidata serves as an important source for the extraction and processing of named entities (Chen et al., 2022c). We computed the proportion of extracted entities that could not be matched to any valid QID. As shown in Table 10, the average proportion of missing QIDs ranges from approximately 26% to 35% across different topics. We manually reviewed the entities that could not be matched to any valid QID and found clear signs of hallucination. In particular, the model frequently generated non-existent items, such as fictional book titles, songs, or dishes. This phenomenon poses a challenge for future efforts to manually complete missing QIDs, as hallucinated entities cannot be trivially verified or linked to real-world knowledge bases.

Topic	Prompt
book	Extract book titles, book genres, author names, places, or reader names. For example: {"<book_title>": "book_title", "<author_name>": "author_name"}.
food	Extract dish names, dish categories, ingredient categories, specific ingredients, places, or person names. For example: {"<dish_name>": "dish_name", "<specific_ingredient>": "specific_ingredient"}.
music	Extract song names, music genres, artist names, places, or listener names. For example: {"<song_name>": "song_name", "<artist_name>": "artist_name"}.
clothing	Extract clothing types, fashion styles, places, or person names. For example: {"<clothing_type>": "clothing_type", "<fashion_style>": "fashion_style"}.
transportation	Extract modes of transport, vehicle types, places, or person names. For example: {"<mode_of_transport>": "mode_of_transport", "<vehicle_type>": "vehicle_type"}.
beverage	Extract drink names, beverage categories, places, or person names. For example: {"<drink_name>": "drink_name", "<beverage_category>": "beverage_category"}.

Table 5: Prompts used for entity extraction per topic.

Language	Model	Model Output (Original)	English Translation
Hausa	Mistral	Wannan abinci yana cikin gida sosai amma...	This dish is very common in the home but... (sentence incomplete)
Hausa	LLaMA	Tsarin ya da kullum bai dace ba da...	Today's system doesn't match with... (unrelated to prompt)
Yoruba	Mistral	Itàn àwón ènìyàn tí a kò m...	A story of people we do not know... (off-topic narrative)
Yoruba	LLaMA	Bni, mo f kí o jun...	Yes, I want you to eat... (no connection to the cultural prompt)

Table 6: Examples of poor model outputs in Hausa and Yoruba, along with manual English translations.

Model Output (Yoruba)	English Translation
Nàìjíríà ní oríl-èdè tí w n ò jún al lóní nítòrí àrùn HIV.	Nigeria is a country where people dine tonight due to HIV.
Blù Adétkunmi. Blù Adétkunmi m tí Nàìjíríà.	Football: Adétkunmi. Adétkunmi is a child from Nigeria.
*K itàn kékeré... s àwn ni yóò s àwn nínú àwn yìn...'	*Write a short story... made people gossip behind closed doors.'

Table 7: Examples of incoherent or culturally inappropriate model outputs in Yoruba.

C.3 Entity Statistics

Table 11 presents the number of unique entities (QIDs) and original labels extracted per topic. The resolve rate, computed as the ratio of unique QIDs to labels, reflects how effectively multilingual surface forms are semantically aligned to the same underlying entity. A lower resolve rate suggests better cross-lingual linking, as more diverse expressions are correctly mapped to shared concepts.

Lang.	Precision (%)	Recall (%)	F1 Score (%)
Arabic	86.30	97.00	91.23
Chinese	88.00	99.00	93.56
English	90.50	99.70	94.96
German	94.52	98.57	96.50
Hindi	87.90	97.10	92.14
Italian	86.13	98.33	91.83
Japanese	89.02	97.33	92.99
Korean	88.11	96.92	92.31
Persian	83.72	98.18	90.38
Spanish	90.42	97.42	93.79
Thai	89.94	98.62	94.08
Traditional Chinese	87.15	97.50	92.04
Turkish	89.71	97.60	93.49

Table 8: NER performance (Precision, Recall, F1) of ChatGPT on different languages. Recall is generally higher, indicating few missing entities, while precision is lower due to over-extraction.

Domain	F1 (%)
Food	91.7
Beverages	91.4
Clothing	96.1
Books	97.4
Music	96.7
Transportation	97.0

Table 9: NER F1 scores across domains in English.

D Results of Granularity

Figure 9 shows presents radar charts of seven models, where each plot shows the average granularity across all topics per language. Each colored line

Topic	Average Proportion of Missing QIDs
Clothing	35.92%
Book	33.25%
Beverage	30.59%
Transportation	25.98%
Music	31.11%
Food	32.18%

Table 10: Average proportion of missing QIDs per topic

Topic	Unique QIDs	Unique Labels	Resolve Rate
beverage	1,152	3,644	12.6%
book	6,978	12,710	12.5%
clothing	2,316	7,432	6.9%
food	6,452	19,202	9.7%
music	6,158	9,964	13.8%
transportation	986	2,644	20.0%

Table 11: Label-to-entity resolution statistics across topics. Each label is linked to a Wikidata QID, enabling cross-lingual semantic alignment.

represents a different country or region mentioned in the prompt. While some countries may appear as outliers, we observe that the overall granularity for a given model-language pair remains relatively stable, may suggesting that the granularity of entity generated by models is invariant to the country or region referenced in the prompt.

E Results of Diversity of Native and Non-native Language

Table 12 reports the average diversity (with variance) of each model across all (*language, topic, country*) combinations. Notably, **DeepSeek** exhibits the lowest overall diversity (16.53), indicating reduced variation in its outputs. In contrast, models like **Llama3-8B** and **Mistral** show higher diversity levels, suggesting richer cultural awareness.




Figure 8: Number of unique cultural entities (Wikidata QIDs) extracted from model outputs in each language.

Model	Variance (Mean \pm Var)
Aya	25.561 \pm 10.4685
ChatGPT	19.561 \pm 8.9480
DeepSeek	16.526 \pm 3.1470
Llama3_8B	39.263 \pm 18.0389
Llama3_70B	27.605 \pm 17.0474
Mistral	34.246 \pm 3.8520
Qwen	13.254 \pm 3.2906

Table 12: Global model averages diversity with variance across languages

F Results of Culture Specificity

Figures 17–23 illustrate the culture specificity of each model, computed by averaging results across all six topics. In this analysis, culture specificity is defined as the difference in entity representations between culturally contextualized prompts (i.e., those mentioning a specific country or region) and neutral prompts. A higher value indicates that the model adapts its output more strongly in response to the cultural cues embedded in the prompt.

In each figure, black-outlined cells highlight cases where the prompt language and the mentioned country form a native pairing. Interestingly, many of these native language–country pairs already exhibit relatively high specificity under the neutral prompt condition, suggesting that language itself can act as an implicit cultural signal, even in the absence of explicit country references.

Across all models, a consistent trend emerges: English exhibits the highest specificity gain, indicating that models tend to adjust their outputs most when prompts are in English. This may reflect English’s central role in pretraining corpora and its exposure to culturally diverse contexts. Additionally, as shown in Figure 23, Chinese achieves a high specificity gain in the Qwen model, potentially reflecting stronger native-language adaptation capabilities.

G Results of Culture Consensus Results

To measure cultural consensus, we compute the pairwise overlap rate of entity QIDs across different languages when generating responses to semantically equivalent prompts. This overlap reflects the extent to which different languages share similar cultural interpretations of the same prompt.

In Figure 25–30, we present the average consensus scores across all mentioned countries and

six topics. To aid interpretation, we use yellow color coding to highlight the major Europe region. Notably, European languages consistently exhibit higher cultural consensus across models, suggesting a stronger alignment in entity representations within this region. Our findings suggest that cultural signals interact with cross-lingual representations in complex ways, shaping both specificity and consensus across languages. Approaches such as contrastive learning (Chen et al., 2022a) offer a promising direction to more faithfully capturing culture-specific cues while maintaining stronger cross-lingual consensus.

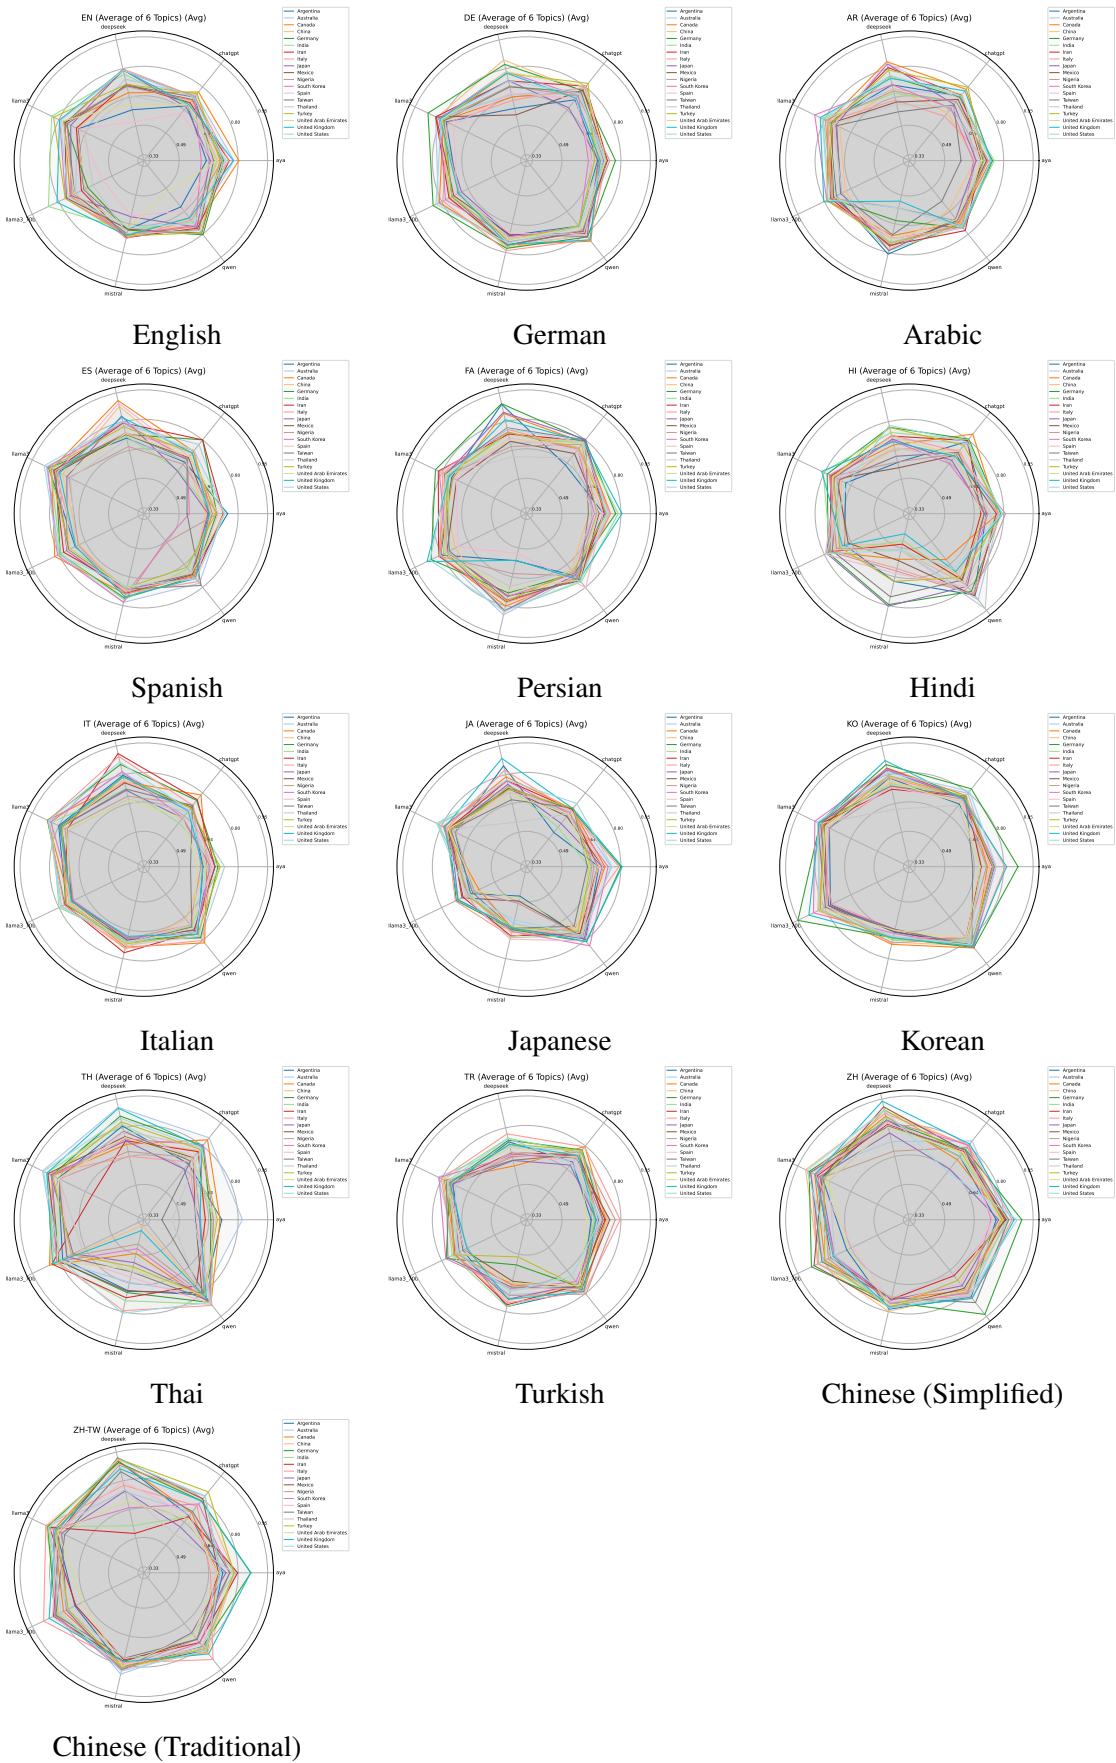


Figure 9: Radar charts of granularity for 15 languages averaged over six cultural topics. Each chart summarizes model-wise generation behavior in the specified language.

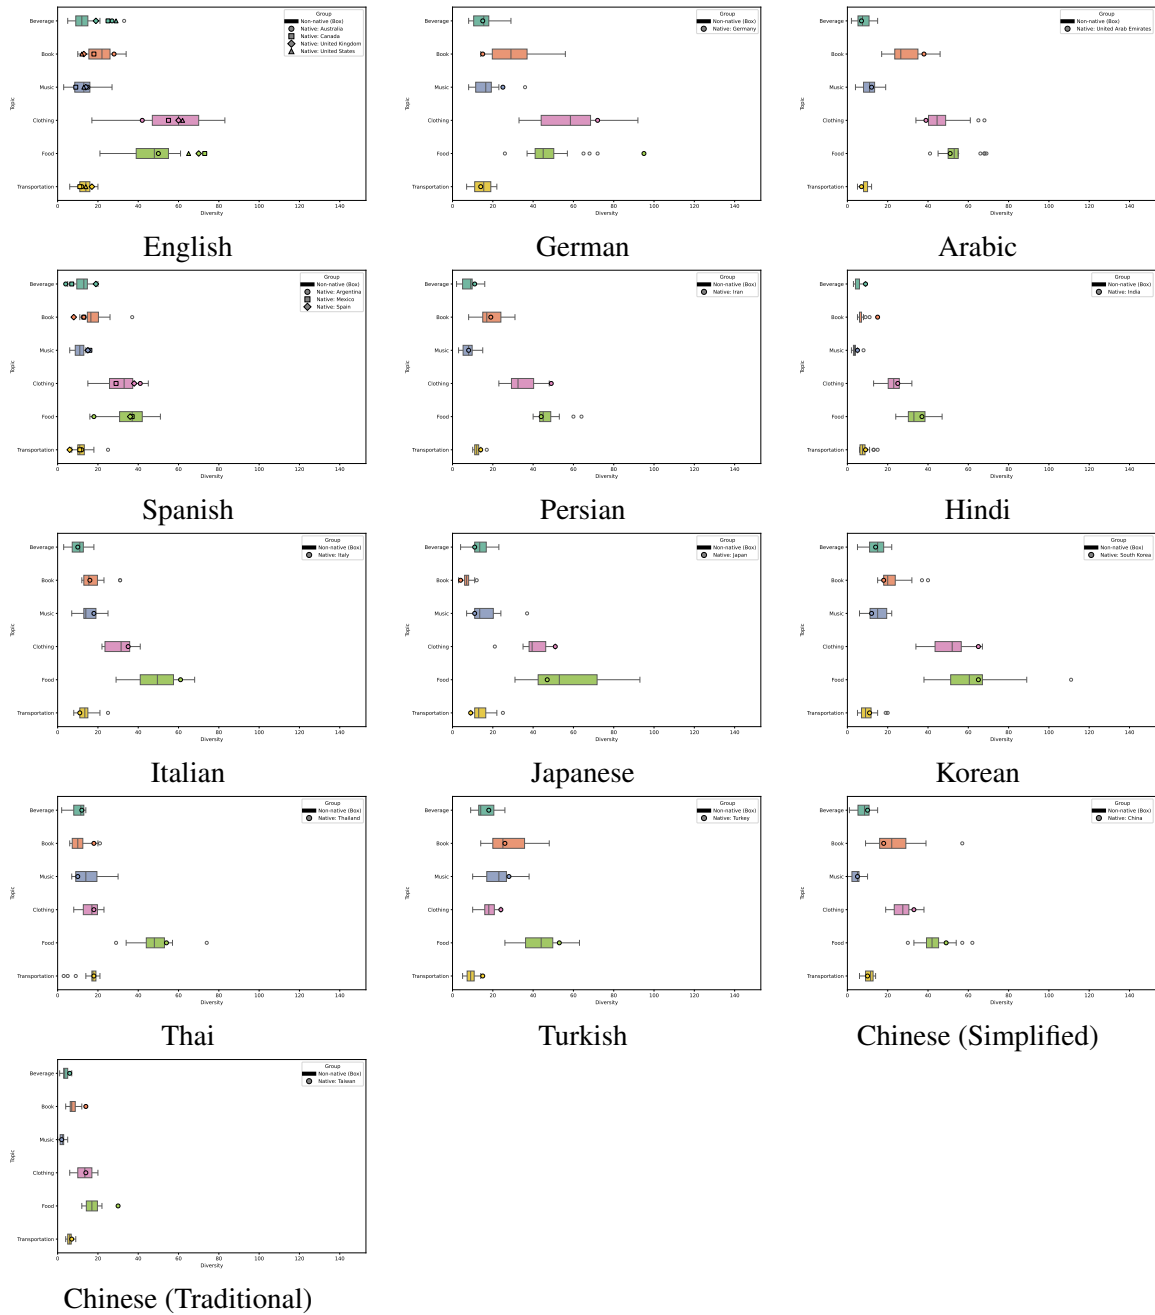


Figure 10: Box plots for diversity comparison between native and non-native languages for CHATGPT across 13 languages.

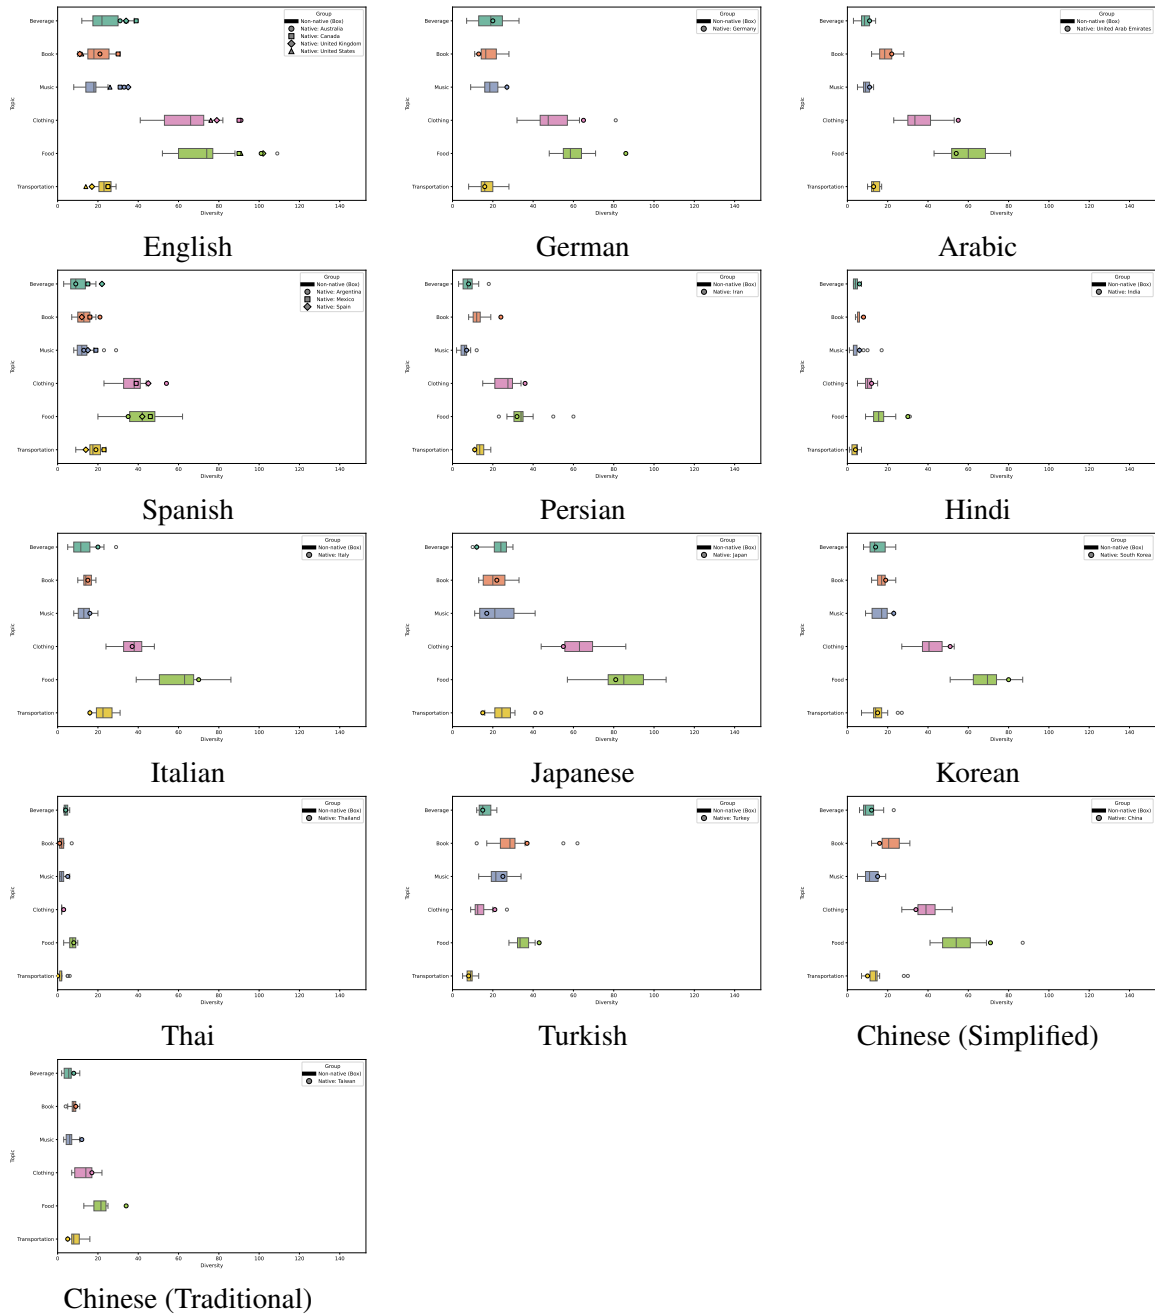


Figure 11: Box plots for diversity comparison between native and non-native languages for AYA across 13 languages.

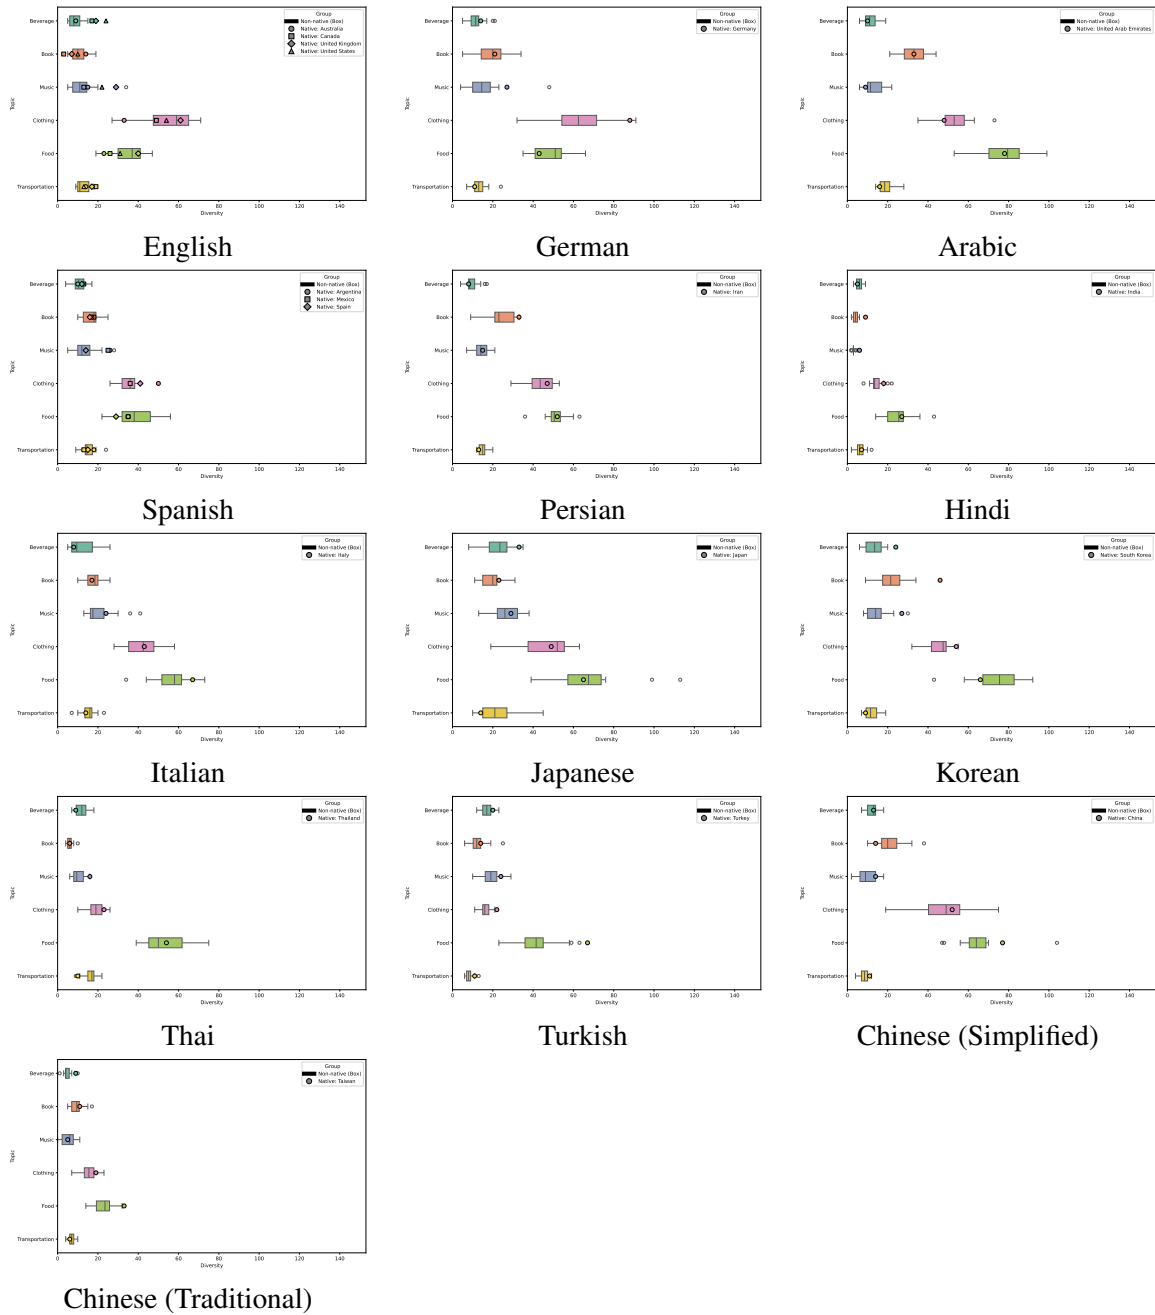


Figure 12: Box plots for diversity comparison between native and non-native languages for LLAMA3-70B across 13 languages.

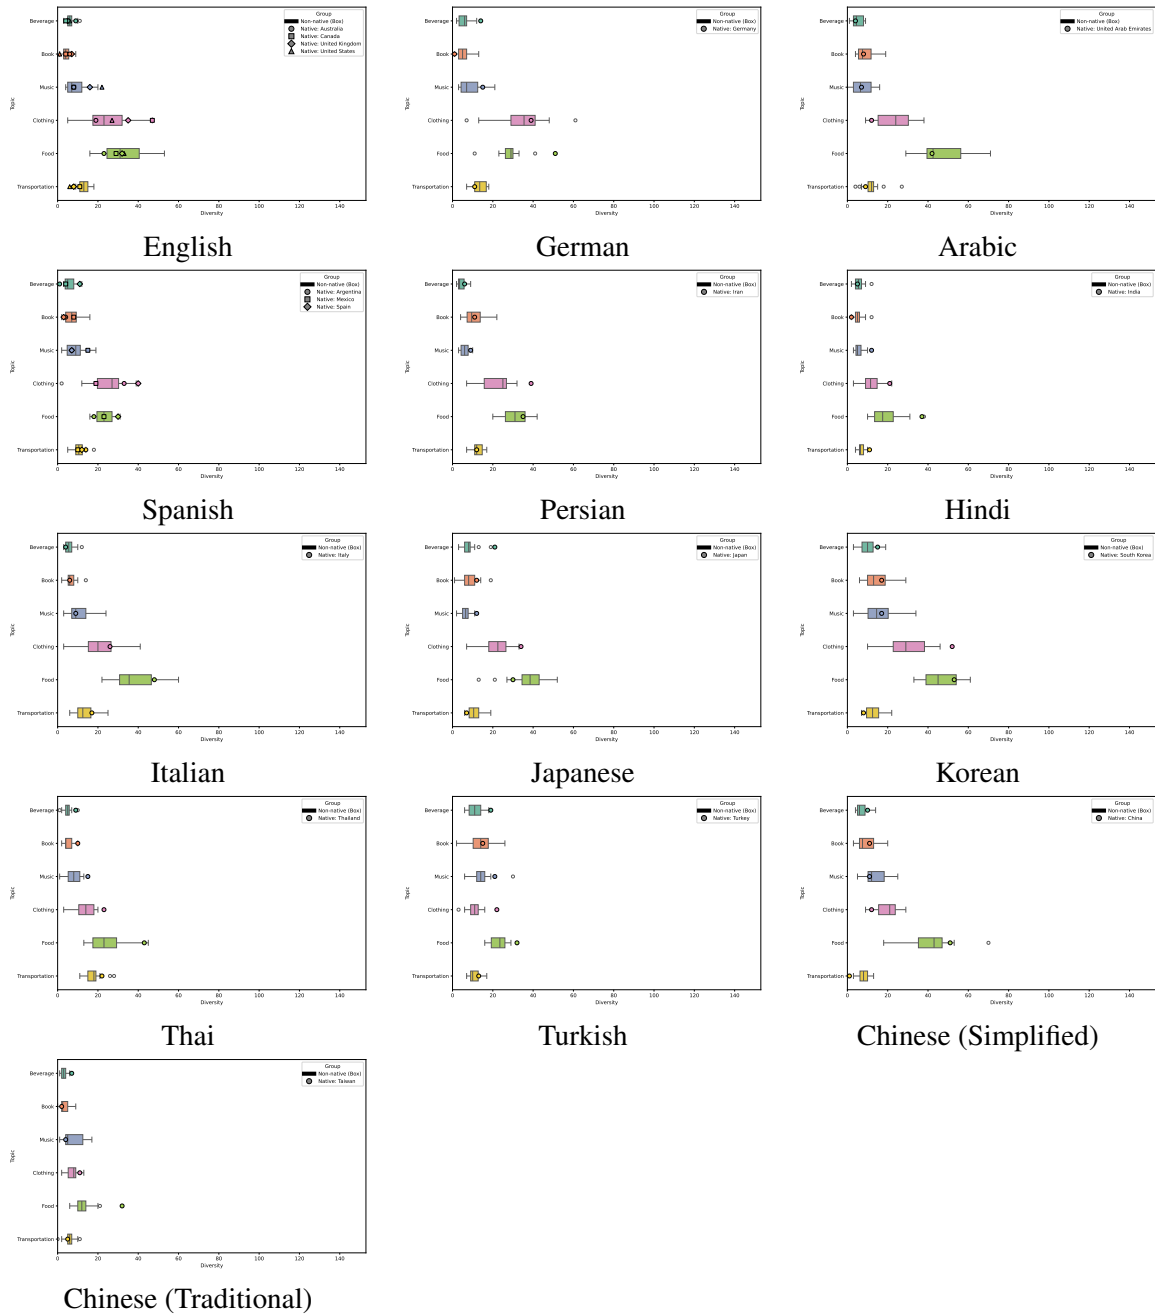


Figure 13: Box plots for diversity comparison between native and non-native languages for DEEPSSEEK across 13 languages.

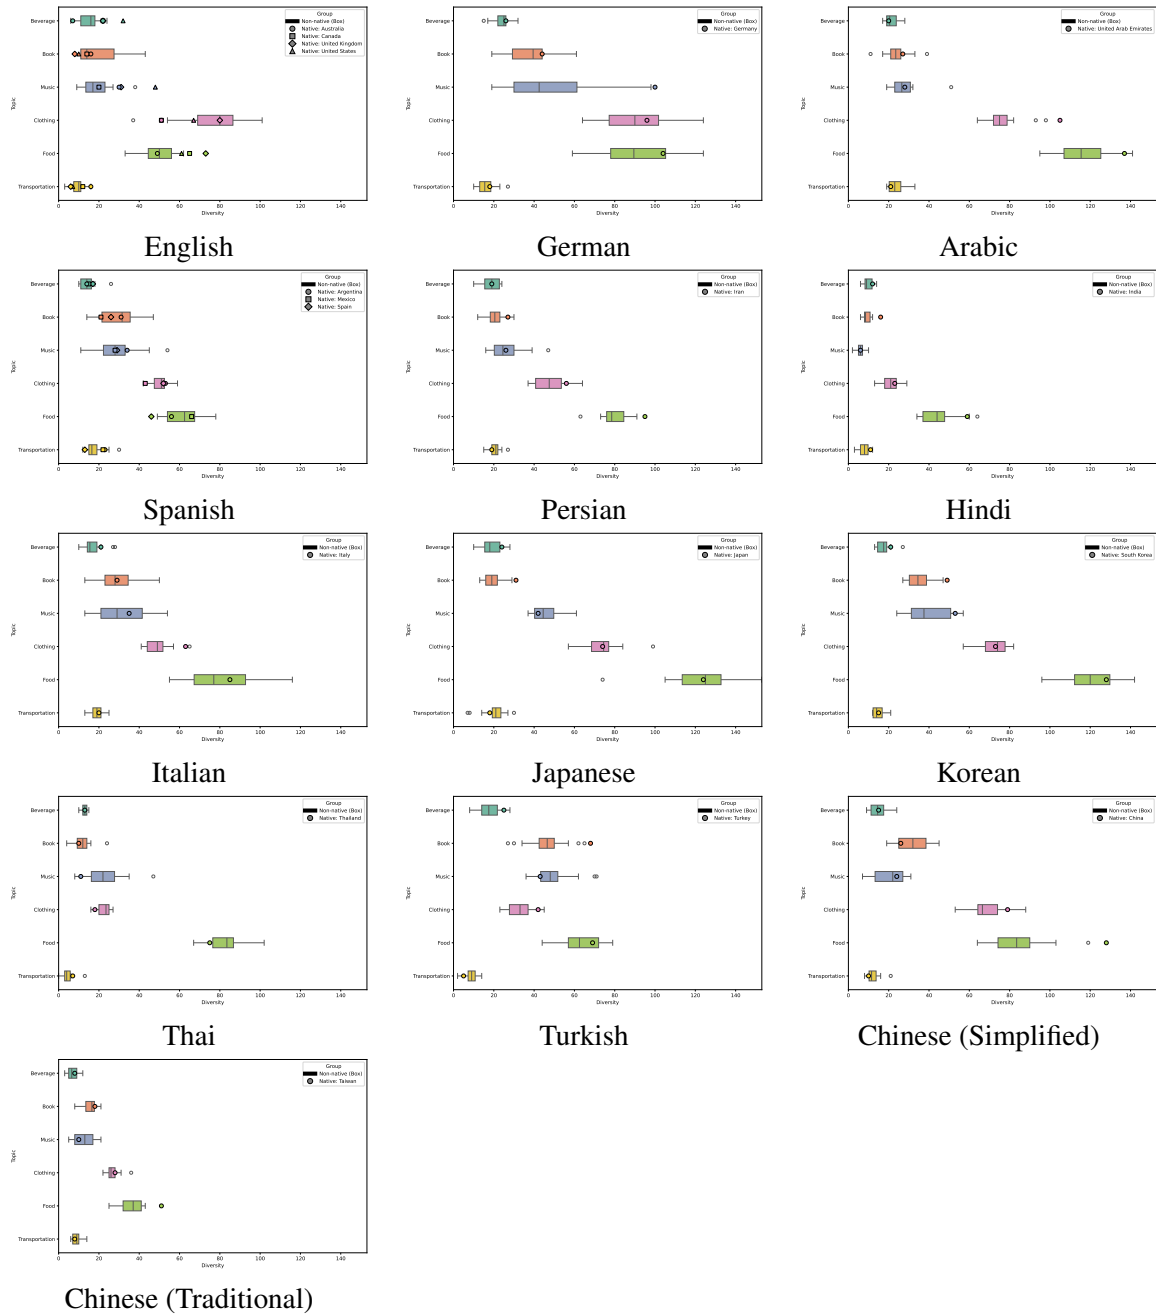


Figure 14: Box plots for diversity comparison between native and non-native languages for LLAMA3-8B

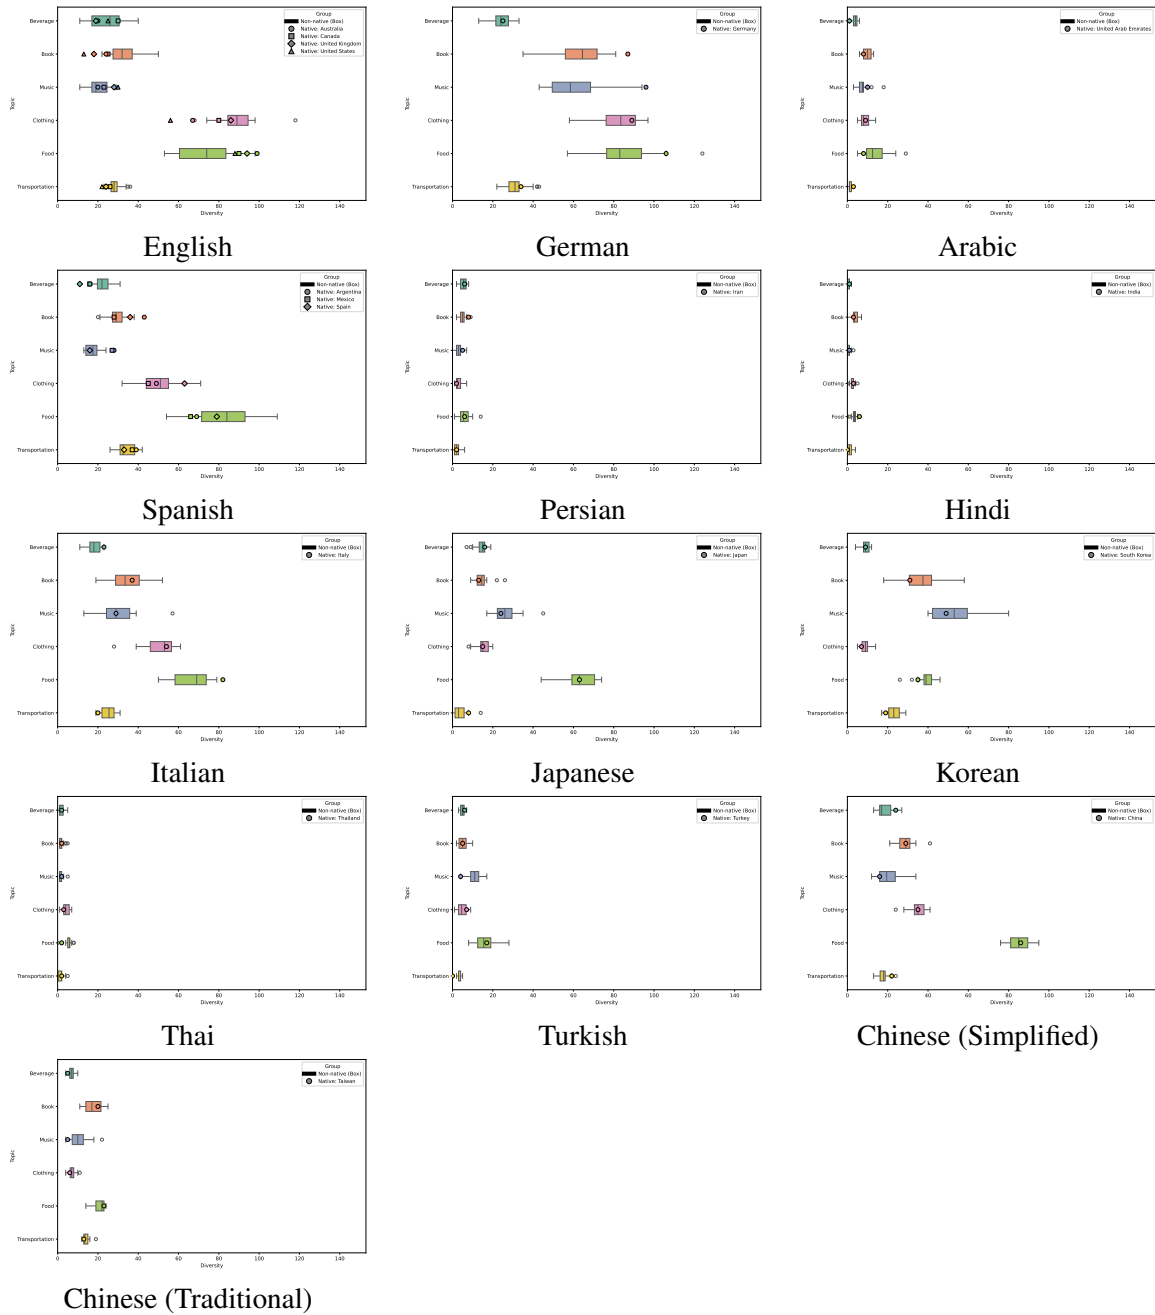


Figure 15: Box plots for diversity comparison between native and non-native languages for MISTRAL across 13 languages.

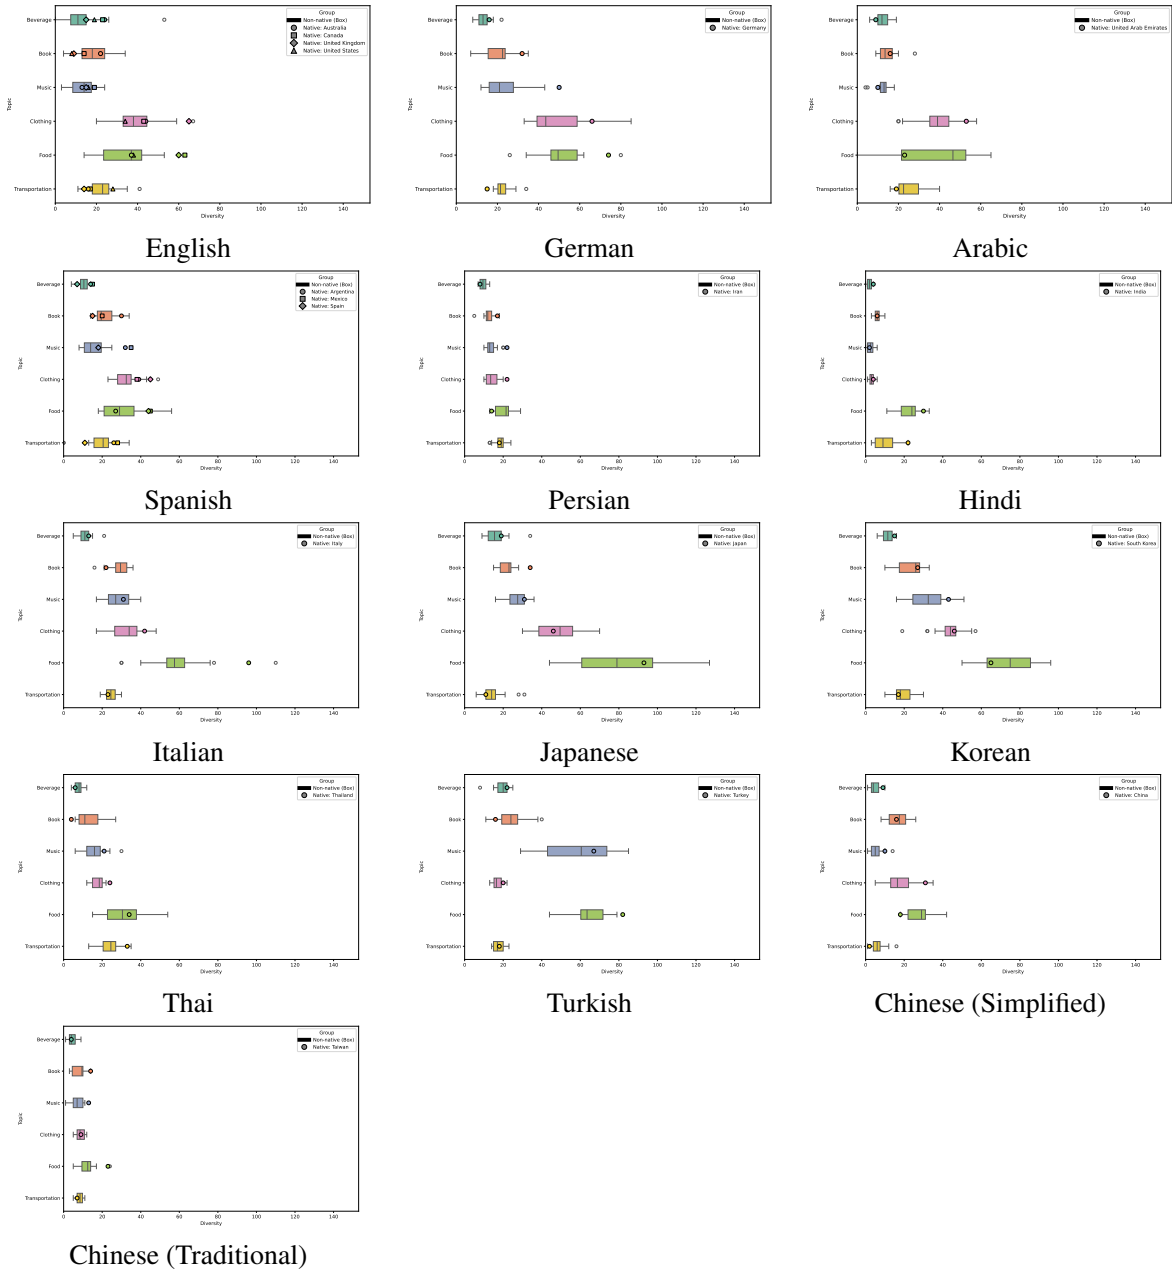


Figure 16: Box plots for diversity comparison between native and non-native languages for QWEN across 13 languages.

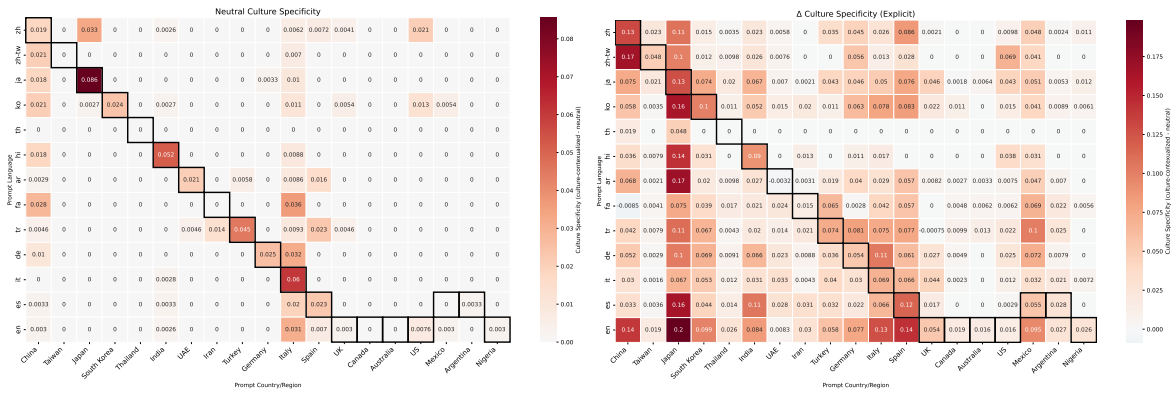


Figure 17: Culture specificity results for **Aya** with a neutral prompt (left) and a prompt explicitly mentioning the country (right, showing delta to neutral).

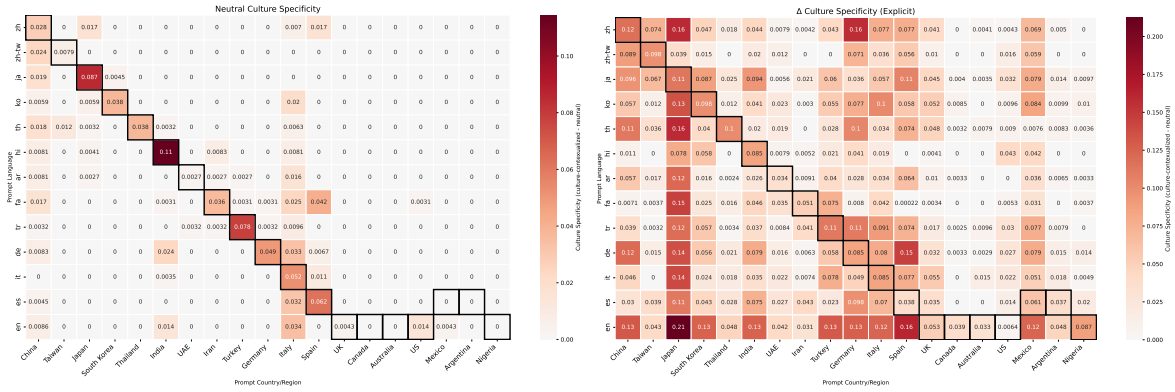


Figure 18: Culture specificity results for **ChatGPT** with a neutral prompt (left) and a prompt explicitly mentioning the country (right, showing delta to neutral).

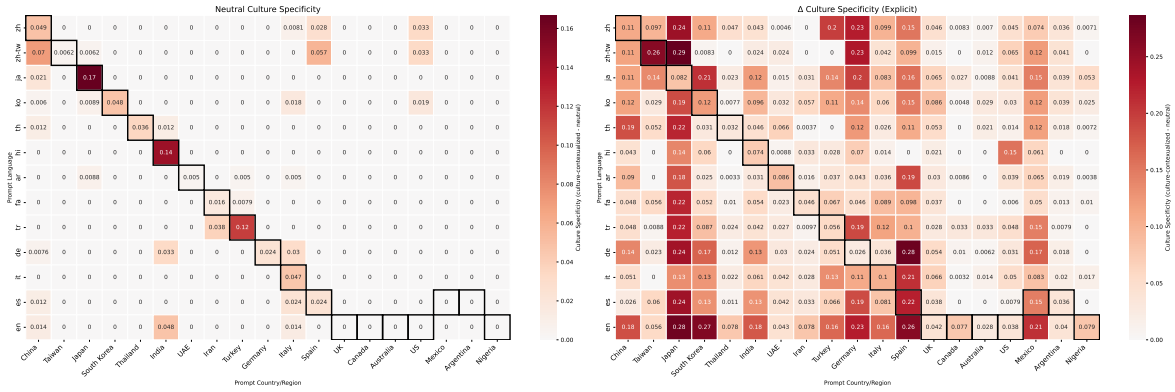


Figure 19: Culture specificity results for **DeepSeek** with a neutral prompt (left) and a prompt explicitly mentioning the country (right, showing delta to neutral).

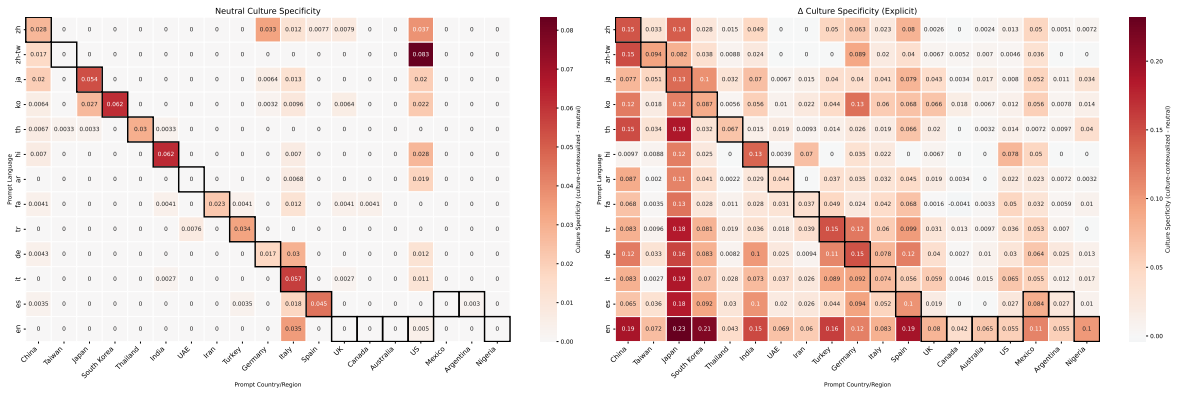


Figure 20: Culture specificity results for **Llama3-70B** with a neutral prompt (left) and a prompt explicitly mentioning the country (right, showing delta to neutral).

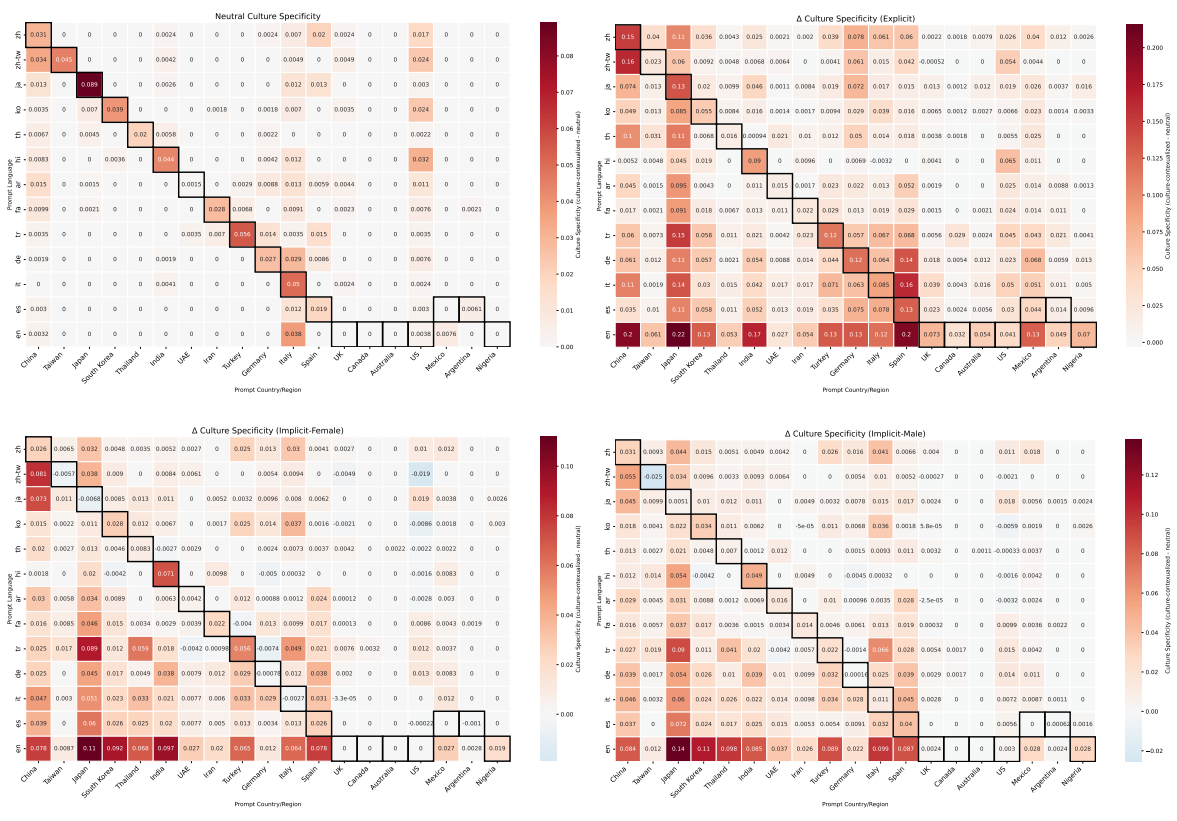


Figure 21: Culture specificity results for **Llama3** with a neutral prompt, a prompt explicitly mentioning the country (showing delta to neutral), and two prompts mentioning a popular female and male name as cultural context (showing delta to neutral).

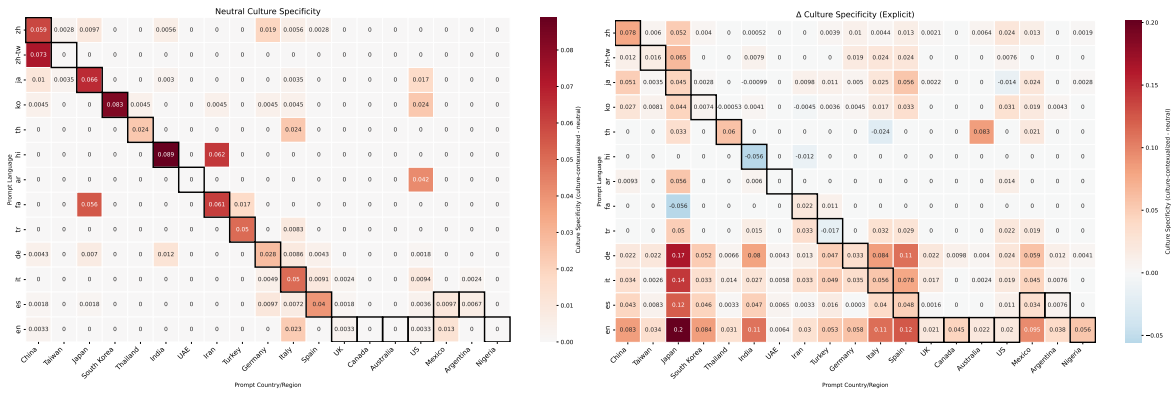


Figure 22: Culture specificity results for **Mistral** with a neutral prompt (left) and a prompt explicitly mentioning the country (right, showing delta to neutral).

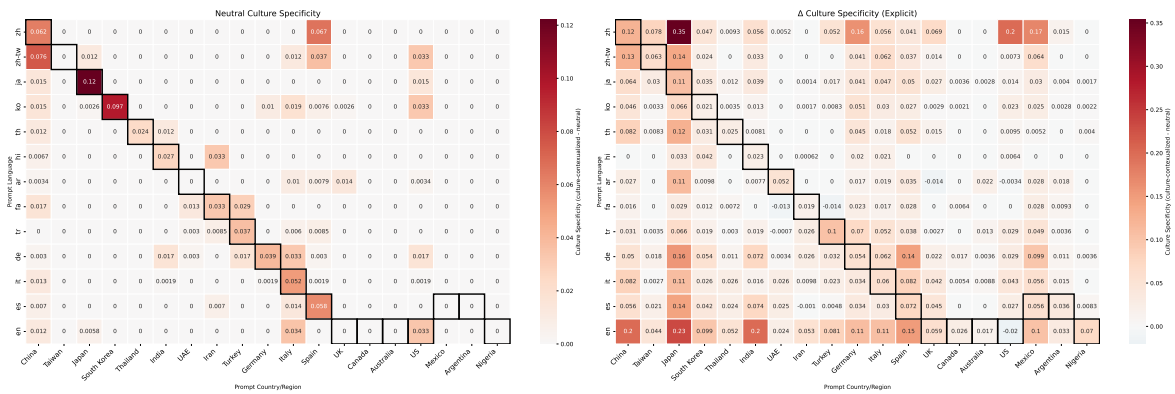


Figure 23: Culture specificity results for **Qwen** with a neutral prompt (left) and a prompt explicitly mentioning the country (right, showing delta to neutral).

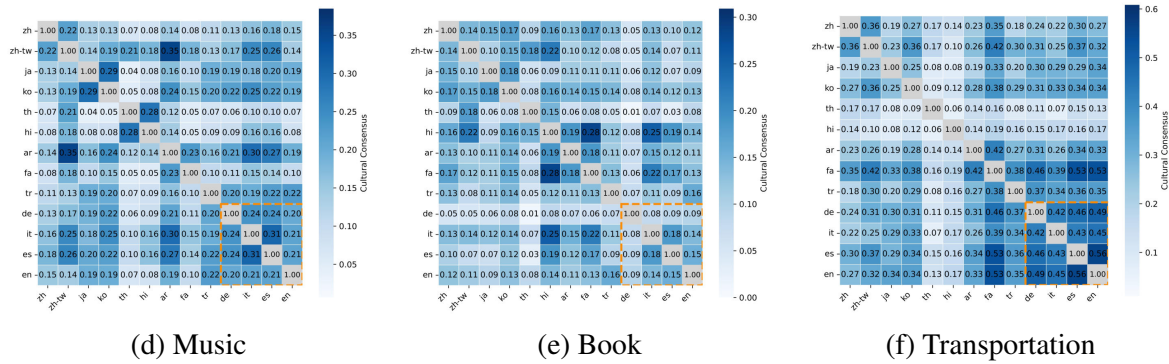
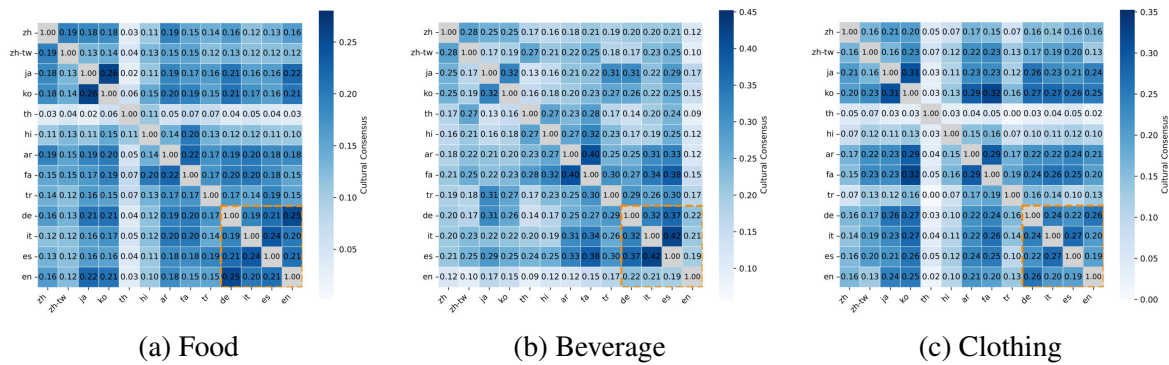


Figure 24: Culture consensus results for **Aya** across six topics. Yellow colors are used to distinguish Europe language region.

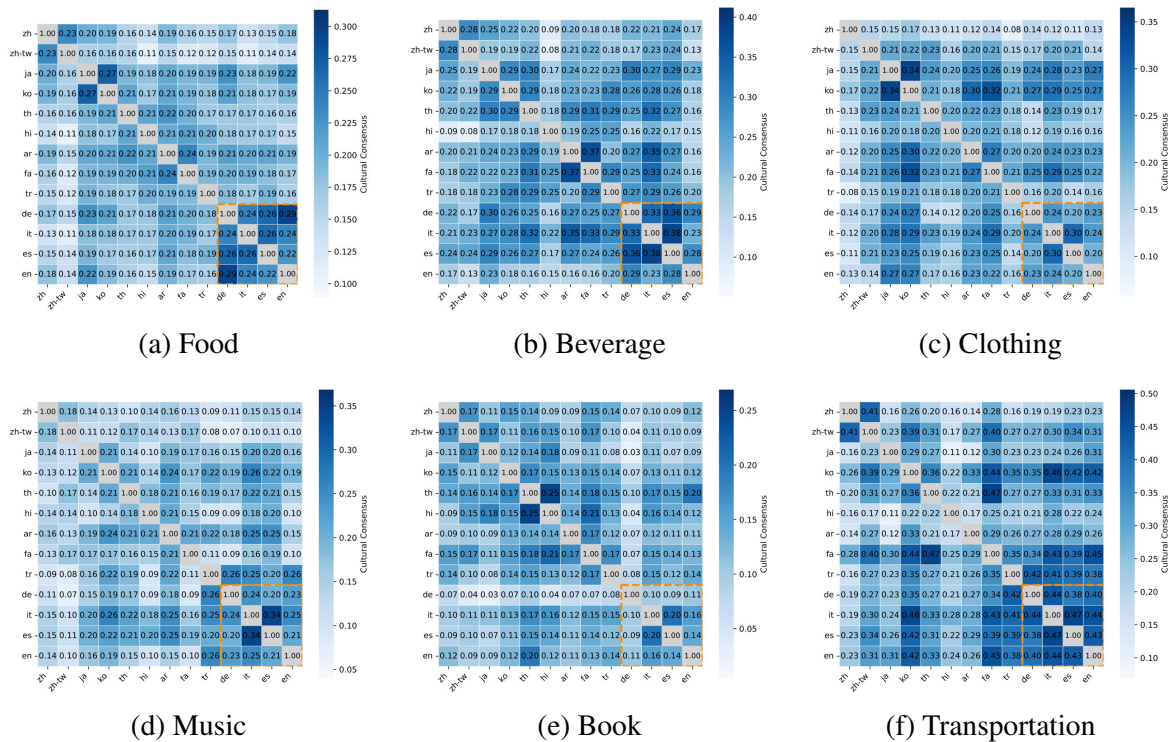


Figure 25: Culture consensus results for **ChatGPT** across six topics. Yellow colors are used to distinguish Europe language region.

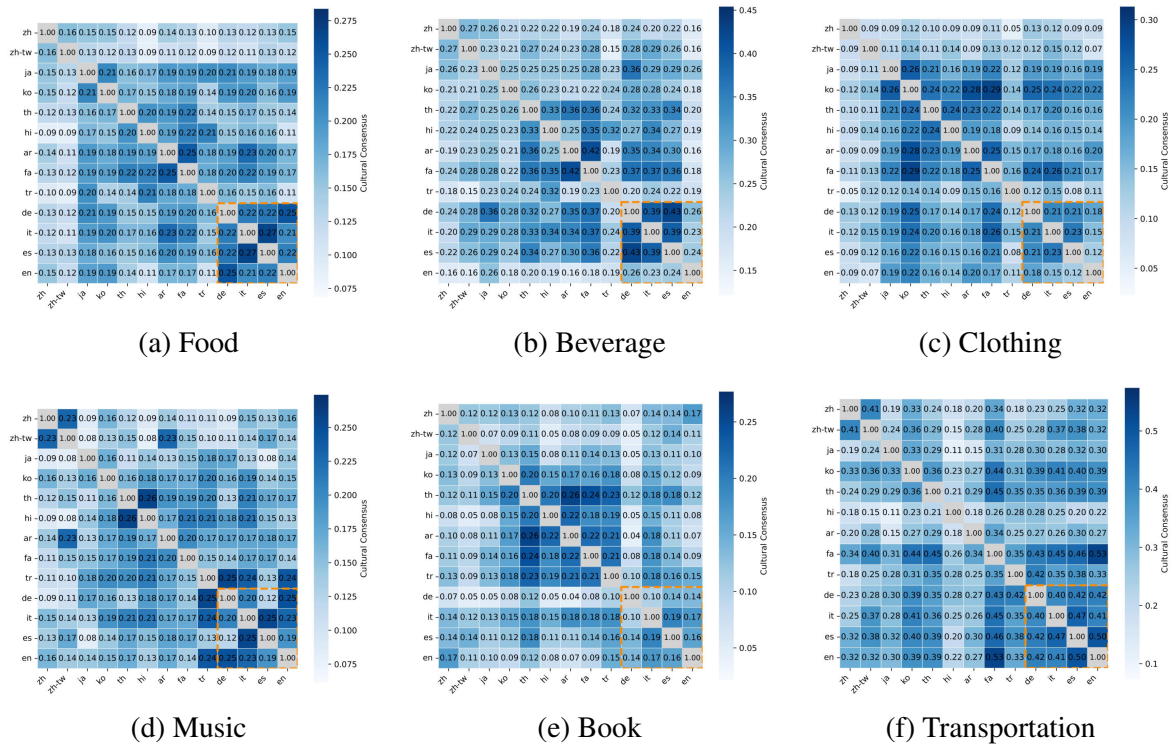


Figure 26: Culture consensus results for **DeepSeek** across six topics. Yellow colors are used to distinguish Europe language region.

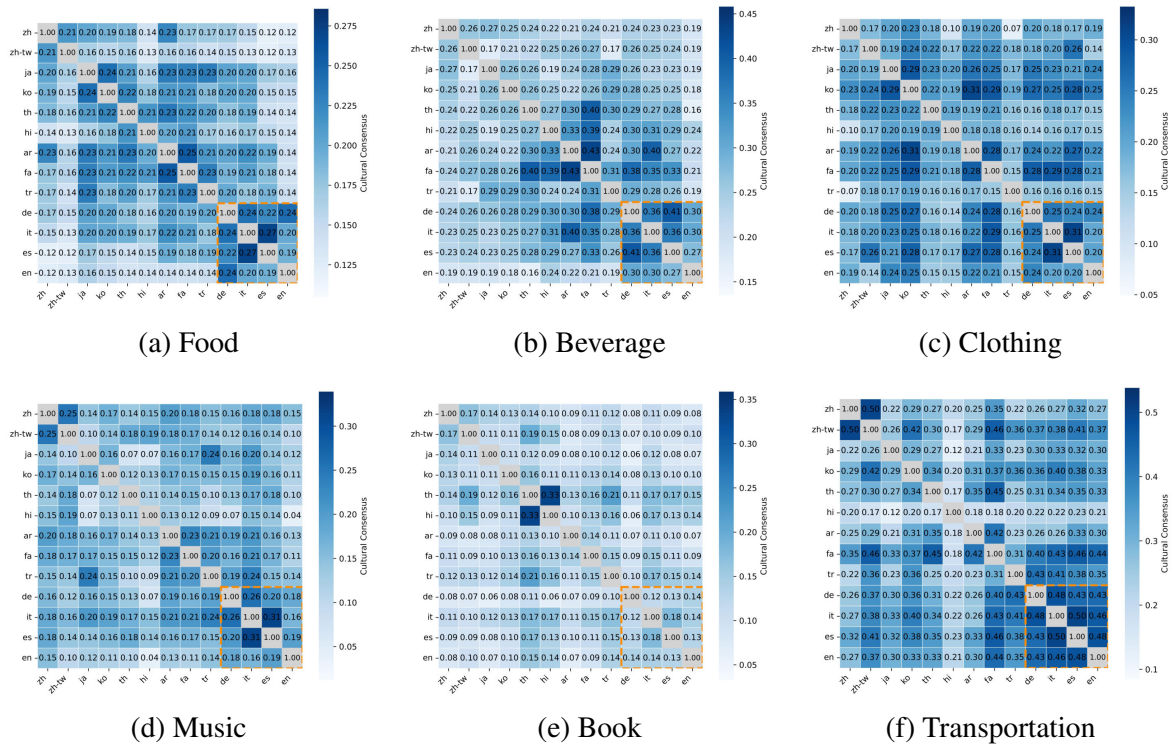


Figure 27: Culture consensus results for **LLaMA3-70B** across six topics. Yellow colors are used to distinguish Europe language region.

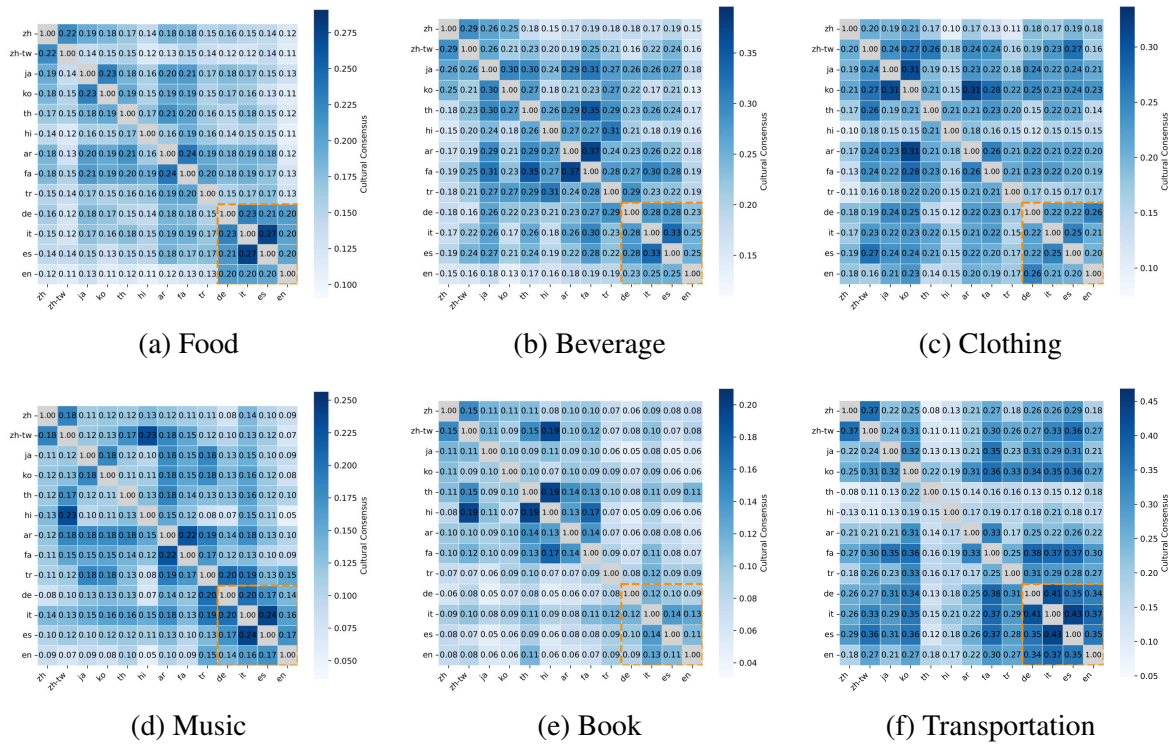


Figure 28: Culture consensus results for **LLaMA3** across six topics. Yellow colors are used to distinguish Europe language region.

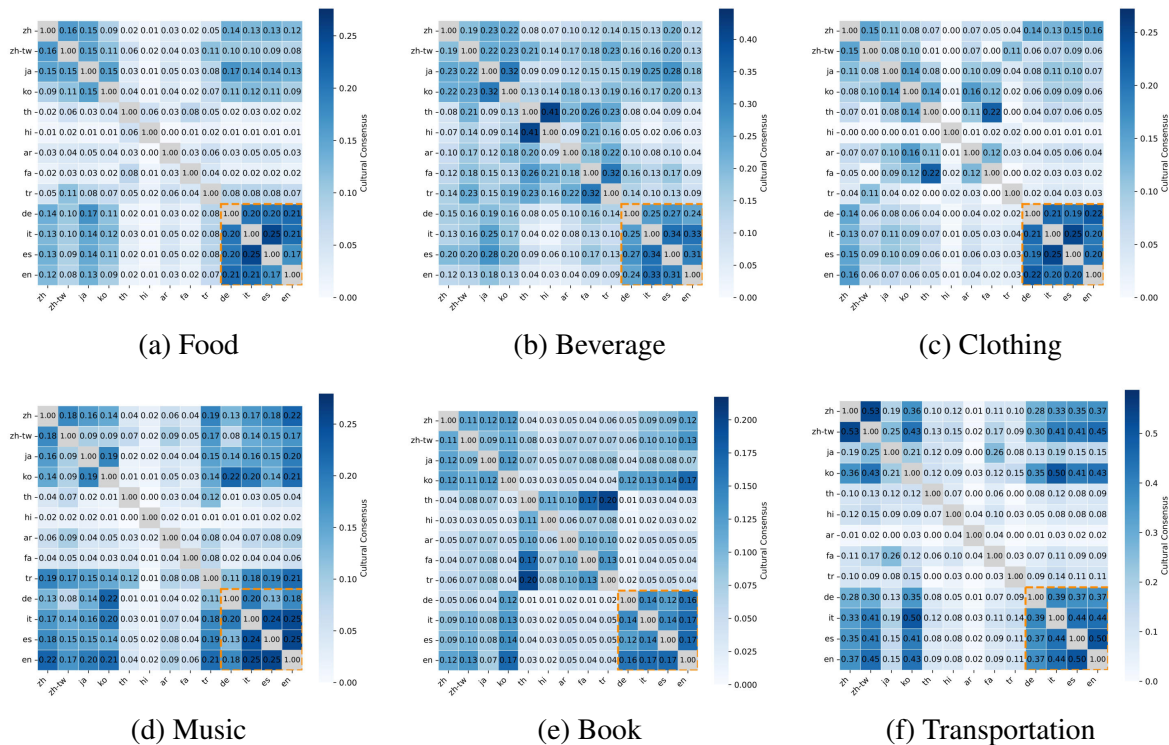


Figure 29: Culture consensus results for **Mistral** across six topics. Yellow colors are used to distinguish Europe language region.

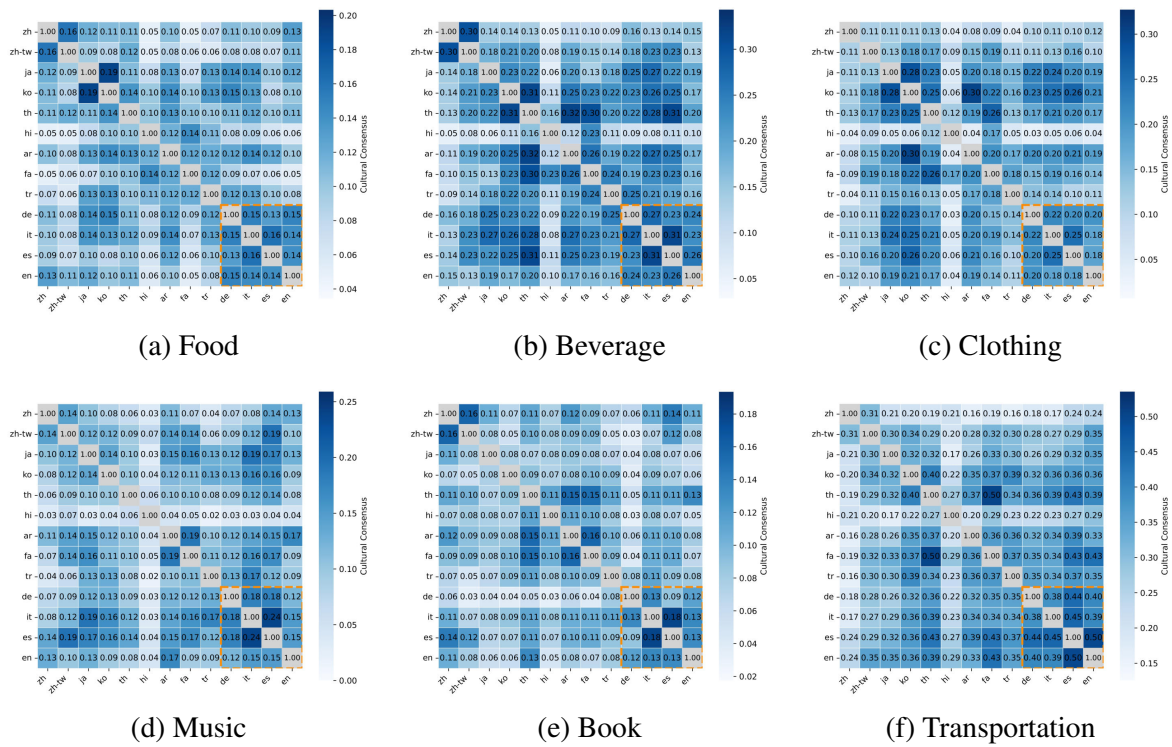


Figure 30: Culture consensus results for Qwen across six topics. Yellow colors are used to distinguish Europe language region.