

# Towards Geo-Culturally Grounded LLM Generations

Piyawat Lertvittayakumjorn<sup>\*†</sup>, David Kinney<sup>\*†‡</sup>,  
Vinodkumar Prabhakaran<sup>†</sup>, Donald Martin, Jr.<sup>†</sup>, Sunipa Dev<sup>†</sup>

<sup>†</sup>Google   <sup>‡</sup>Washington University in St. Louis  
{piyawat, vinodkpg, dxm, sunipadev}@google.com, kinney@wustl.edu

## Abstract

Generative large language models (LLMs) have demonstrated gaps in diverse cultural awareness across the globe. We investigate the effect of retrieval augmented generation and search-grounding techniques on LLMs' ability to display familiarity with various national cultures. Specifically, we compare the performance of standard LLMs, LLMs augmented with retrievals from a bespoke knowledge base (i.e., KB grounding), and LLMs augmented with retrievals from a web search (i.e., search grounding) on multiple cultural awareness benchmarks. We find that search grounding significantly improves the LLM performance on multiple-choice benchmarks that test propositional knowledge (e.g., cultural norms, artifacts, and institutions), while KB grounding's effectiveness is limited by inadequate knowledge base coverage and a sub-optimal retriever. However, search grounding also increases the risk of stereotypical judgments by language models and fails to improve evaluators' judgments of cultural familiarity in a human evaluation with adequate statistical power. These results highlight the distinction between propositional cultural knowledge and open-ended cultural fluency when it comes to evaluating LLMs' cultural awareness.

## 1 Introduction

Contemporary large language models (LLMs) are pre-trained on huge corpora of natural language text (Radford et al., 2019) and then fine-tuned using human feedback to improve their quality (Bai et al., 2022). During both processes, it is possible for text from a particular culture or cultures to be over-represented in the training data (Dodge et al., 2021) and for the perspectives, norms, and mores of specific cultures to be over-represented in the feedback from human evaluators (Prabhakaran et al., 2022; Atari et al., 2023). Consequently, there is growing recognition of generative LLMs' shortcomings in representing and serving people from diverse geo-cultural backgrounds at the global scale (Adilazuarda et al., 2024; Pawar et al.,

2024; Agarwal et al., 2025). Models tend to stereotype different cultures (Jha et al., 2023; Bhutani et al., 2024), erase and simplify their representation (Qadri et al., 2025), and provide very limited knowledge and context about artifacts and norms that are salient to them (Myung et al., 2024; Rao et al., 2024). Despite these gaps, strategies for eliciting culturally appropriate content from the models remain under-explored, with only some investigation of prompt engineering (Rao et al., 2023; Wang et al., 2024) and model fine-tuning (Chan et al., 2023; Li et al., 2024a,b) if not pre-training on more diverse non-English data.

Here, we study two strategies to improve cultural awareness<sup>1</sup> of LLM generation using external knowledge. In the first strategy, we construct a bespoke cultural knowledge base (KB) and apply a retrieval augmented generation (RAG) technique (Lewis et al., 2020; Gao et al., 2023) so that the input to the language model includes relevant cultural text from the knowledge base for better generation. In the second strategy, we use a commercially-available search-grounding generation API which translates user prompts into a web search query, uses it to retrieve relevant pieces of text from the Internet, and grounds the LLM generation on the retrieved text. We call these two strategies **KB-grounding** and **search-grounding**, respectively.

In our experiments, we highlight the necessity of a multi-pronged approach to evaluating cultural awareness in language model generations. Specifically, we leveraged multiple benchmarks (Myung et al., 2024; Rao et al., 2024; Bhutani et al., 2024) to evaluate cultural knowledge and the ability to avoid cultural stereotyping of the models equipped with the two strategies, compared with the vanilla generation baseline. We also conducted a human evaluation of open-ended model responses to various prompts designed to test cultural fluency, wherein evaluators from a specific national culture rated how well the model's output reflected a culturally familiar perspective. The results from both experiments shed light on the pros and cons of the two strategies. Finally, we conclude the paper by discussing key findings and offering suggestions for future work.

<sup>1</sup>We use the term "cultural awareness" to refer to the general ability of LLMs to work well across cultures. It is conceptually similar to other terms in the literature, e.g., "cultural appropriateness" and "cultural informedness." However, we are not attempting to define them precisely in this paper.

\* equal contribution

| Source       | Description               | # Docs  |
|--------------|---------------------------|---------|
| CultureAtlas | Wikipedia text            | 239,376 |
| Cube         | Artefact names            | 198,896 |
| CultureBank  | Situation-based practices | 22,990  |
| SeeGULL      | Stereotypes               | 6,871   |

Table 1: Sources of documents in our cultural KB.

## 2 Improving Cultural Awareness by Retrieving External Knowledge

Retrieval augmented generation (RAG) is a technique for enhancing the quality of large language model generation. To implement RAG, a user prompt is first used to retrieve relevant information (e.g., text or documents) from a database, which is then added to the prompt before being passed to a generative LLM to produce a grounded response (Lewis et al., 2020; Gao et al., 2023). RAG has shown to be effective in several applications, particularly those involving tasks that the base LLM was not well-trained for such as fact verification (Asai et al., 2023; Singal et al., 2024; Khaliq et al., 2024) and domain-specific question answering (QA) (Seo et al., 2024; Xiong et al., 2024; Kim and Min, 2024). Some existing work also creates bespoke knowledge bases for use with RAG to tailor outputs to their specific applications (Sun et al., 2025; Li et al., 2024c). Meanwhile, instead of querying a KB, other techniques retrieve external information by searching the internet, which allows access to up-to-date knowledge with high-quality ranking outputs (Fan et al., 2024; Shuster et al., 2022; Lazaridou et al., 2022; Yao et al., 2022; Nakano et al., 2021; Komeili et al., 2022). In light of these successes, we study both KB-grounding and search-grounding as techniques for improving the cultural awareness of LLM generations.

### 2.1 The Knowledge Base Grounding Strategy

We compiled culturally salient data from four large sources—CultureAtlas (Fung et al., 2024), Cube (Kannen et al., 2024), CultureBank (Shi et al., 2024), and SeeGULL (Jha et al., 2023)—to serve as our knowledge base as listed in Table 1. For each entry in each source, we converted it into text (if it was not already), embedded it into a vector, and added it to a vector store for querying.

Fig. 1 (top) shows how the KB-grounding strategy works. First, a query rewriter extracts the important parts of an incoming prompt to form a query. We can configure this step to test different KB queries (e.g., whether to include choices of a multiple-choice question prompt in the query). Next, we use the query to retrieve  $n$  documents from the KB. We optionally check whether each of the documents is indeed relevant to the original prompt by prompting the base LLM and include only the  $k$  relevant ones in the prompt. We call the process with the relevancy check step *selective RAG*, while *non-selective RAG* includes all  $n$  doc-

uments in the prompt. Then, we feed the augmented prompt to the LLM to get a raw answer. For multiple-choice question answering (QA) tasks, the LLM sometimes does not strictly follow the formatting instruction in the prompt, resulting in various surface forms of the same answer. For example, the raw outputs for the answer “1) Yes” we found include, e.g., “1”, “1)”, “Yes”, “1) Yes”, “Answer: Yes”, “Answer: 1) Yes”, “\*\*Answer\*\*”: Yes”, etc., sometimes with additional explanations. Therefore, we create a method, called a manual verbalizer, to normalize raw answers and map them to one of the choices so we can compute the model accuracy (similar to verbalizers in LLM-based text classification (Schick and Schütze, 2021; Thaminkaew et al., 2024)). By contrast, we use the raw answer as the final output for an open-ended generation task.

For implementation details, we created our vector store on Google Vertex AI using textembedding-gecko@003 as the embedding model. In the experiments, we retrieved  $n = 5$  most similar texts from the vector store; however, the number of texts actually used ( $k$ ) could be lower than 5 for the selective RAG approach. In Section 3, we applied KB-grounding to three LLMs: gemini-flash-1.5 (Gemini) (Team et al., 2024), gpt-4o-mini (GPT) (OpenAI et al., 2024) and olmo2-7b (OLMo) (OLMo et al., 2025), all with the temperature of 0.5. See the Appendix for details about the KB, the queries, and the prompt templates for relevancy check and answer generation.

### 2.2 The Search Grounding Strategy

Fig 1 (bottom) outlines how the search-grounding strategy works. We feed the original prompt to the search-grounding generation API which converts the prompt into a search query, inputs it to a search engine, and obtains relevant text from the web pages returned by the search. Thus, the search-grounding strategy retrieves text that is relevant to the prompt by effectively exploiting proprietary page-ranking algorithms used in contemporary search engines to identify web pages that are relevant to the prompt, and then extracting and checking the relevancy of specific text from those pages using text extraction and relevancy-checking techniques that are not, currently, publicly available. The retrieved relevant text is then integrated into the prompt, at which point the LLM generates a response based on this augmented prompt. Thus, the search-grounding strategy effectively replaces the bespoke KB in Section 2.1 with the entire web, and replaces the vector-based KB querying with retrieving prompt-relevant content with a powerful search engine. We implemented search grounding on Google Vertex AI. This enables retrieval of prompt-relevant content from the Google search engine, which is then integrated into the prompt that is given to the Gemini LM to generate a response. Search-grounding is not currently available on APIs for accessing GPT or OLMo, so we only implemented search-grounding for Gemini.




Figure 1: (Top) The knowledge base grounding strategy and (Bottom) The search grounding strategy. Dashed boxes indicate optional steps that are executed only under the annotated conditions. For KB grounding, the same LLM is used for both relevancy check and answer generation.

### 3 Cultural Competence Benchmarks

We evaluated the KB-grounding and search-grounding strategies on multiple-choice cultural QA benchmarks. The key results are discussed in this section, while the full statistical analyses are reported in Appendix A.5.

#### 3.1 Cultural Knowledge

**Setup.** We used two benchmark datasets to test the sensitivity of LLMs with respect to two facets of culture. The first, BLEnD, contains questions about everyday cultural knowledge (such as food, sports, family, and education) in different countries (Myung et al., 2024). We used the  $\sim 24k$  English questions from BLEnD that are related to ten countries<sup>2</sup> represented in our bespoke KB. The second, NORMAD, focuses on cultural norms and values (Rao et al., 2024). Each question is a story in an everyday scenario, and asks whether a character’s action in the story is socially acceptable within the given context (Yes, No, or Neither). We experimented with two types of contexts – *Country* and *Country+Value*. The former specified a country where the story takes place, while the latter additionally indicated the value paradigm the character should adhere to. NORMAD has  $\sim 2.6k$  questions in total.

<sup>2</sup>China, Ethiopia, Greece, Indonesia, Iran, Mexico, South Korea, Spain, the United Kingdom, and the United States.

We compared the vanilla generation with the KB-grounding and the search-grounding generations. For KB-grounding, we considered both the selective and non-selective RAG approaches. For BLEnD specifically, as answer choices are potentially relevant to the question, we considered KB query rewriting both “with choices” and “without choices” included in the query. However, we always included the choices in the prompts for relevancy check and answer generation.

**Results.** Fig. 2 (left) presents the results of the BLEnD benchmark. A repeated-measures ANOVA revealed a significant effect of strategy on accuracy across all three LLMs, meaning that the accuracy of LLM generations differed significantly between at least some of the generation strategies used (i.e., search-grounding, the varieties of KB-grounding, or vanilla). However, the optimal methods differed among the LLMs: search-grounding for Gemini, non-selective KB-grounding (without answer choices in KB queries) for GPT, and both selective KB-grounding (with choices in KB queries) and the vanilla approach for OLMo. Note that although the magnitude of the difference in performance between strategies was relatively small for both Gemini and GPT, the large sample size ( $\approx 24k$  examples) enabled us to detect the statistical significance of these improvements.

We also observed the effectiveness of both KB- and search-grounding in several examples. For instance, the vanilla Gemini incorrectly answered the question, “What is the most popular sport team in Ethiopia? (A) coffee (B) lg twins (C) persepolis (D) real madrid.” It selected (D), possibly due to Real Madrid’s high frequency in the training corpora concerning popular sport teams. In fact, the correct answer is (A), as ‘coffee’ refers to the Ethiopian Coffee Sport Club. This information is available on the internet; hence, it is not surprising that search-grounding correctly answered this question. Notably, search-grounded Gemini achieves 74.2% accuracy for questions related to Ethiopia, significantly higher than 60.3% of the vanilla baseline and 62.9% of the best KB grounding setting. Certain KB-grounding settings also answered the above question correctly, even though ‘Ethiopian Coffee’ was not mentioned in any of the retrieved texts. The only text retrieved by selective RAG KB-grounding was about sports in Ethiopia in general, which possibly reminded the model to avoid answers from other countries. While this worked, it would be better and more reliable if the knowledge base had wider coverage, including direct information about the Ethiopian Coffee team.

Interestingly, the KB grounding strategy did not improve OLMo’s performance on BLEnD. This is likely because OLMo was the weakest model among the three according to its lowest vanilla performance across the benchmarks. It also struggled to utilize the retrieved documents effectively, unlike Gemini and GPT. This issue was prominent when we used the non-selective KB grounding strategy, where the relevancy check step




Figure 2: Performance of all strategies for all models on the BLEnD, NORMAD (Country and Country+Value), and stereotype avoidance benchmarks, with 95% confidence intervals; higher values are better for all plots.

was not applied and the model was presented with all the five retrieved texts, some of which were irrelevant and lengthy. These texts could confuse OLMo and lead it to a wrong answer often because the irrelevant texts were tangentially related to the question or choices but did not contain the correct answer. Furthermore, in the non-selective KB grounding setting, OLMo outputted malformed answers more often. BLEnD requires the answer in the JSON format. Although our manual verbalizer allows some flexibility, it still cannot parse many answers from OLMo with non-selective KB grounding, leading to the lower scores. We hypothesize that the model may have difficulty adhering to the formatting instruction when overloaded with confusing retrieved texts. However, both issues were less pronounced in selective KB grounding, where only relevant texts were included. Consequently, in this case, selective KB grounding (with choices in KB queries) performed comparably to the vanilla baseline.

The middle two charts in Fig. 2 present the results for NORMAD for both Country and Country+Value context types. As with BLEnD, search-grounded Gemini performed best for both context types compared to other strategies. For GPT, selective KB-grounding outperformed alternatives in the Country context, while the vanilla one performed best in the Country+Value context. The trend for OLMo was similar to GPT, except that selective KB-grounding performed comparably to the vanilla for the Country+Value context. We found that KB grounding struggled in this task because, several times, the retrieved texts were related to one aspect of the question but not all. For example, some retrieved texts were about the country but not the exact situation in the question or vice versa. This could cause the LLMs to give wrong answers. The relevancy check step in selective RAG mitigated this issue, albeit to a limited extent. We posit that a better-tuned KB retriever would further address this problem (if the pertinent fact is present in the KB (Joren et al., 2024)).

### 3.2 Stereotype Avoidance

**Setup.** An important aspect of demonstrating cultural awareness is the avoidance of stereotyping behavior.

Hence, we adopted the SeeGULL stereotype avoidance test (Bhutani et al., 2024) to evaluate the studied strategies. Each question asks the model to indicate a person’s nationality based on a stereotype. For example,

“I know a person who is *a fighter*. This person is: (A) Albanian (B) Afghan (C) Filipino (D) Guyanese. Write your answer on the next line, or write (?) if you are unsure.”

In every question, only one country in the choices is stereotypically associated with the question. However, we considered only ‘(?)’ the correct answer, since it reflects the LLM’s ability to avoid stereotypical responses. The benchmark contains 4.6k questions. As with BLEnD, since the choices may be semantically related to the question, we implemented selective and non-selective RAG both with and without answer choices included in the KB query.

Note that our bespoke KB and the resource used to construct this stereotype avoidance benchmark overlap as both rely on SeeGULL (Jha et al., 2023), and this setup is intended. Unlike cultural knowledge evaluation in Section 3.1, in this task, if the model retrieves a SeeGULL stereotype and uses it to answer the question, it is liable to produce an incorrect answer, affirming the stereotypes (rather than answering the correct “unsure”). This setup enables us to understand how well the KB-grounding strategy handles stereotypical facts, which could exist in real-world cultural KBs.

**Results.** Fig. 2 (right) shows that, on this benchmark, Gemini vastly outperformed GPT and OLMo in avoiding stereotyping, with the vanilla Gemini showing the best performance. However, while all other Gemini methods performed relatively well on this task, search-grounding led to a significant degradation in performance, with the model selecting the stereotypical options considerably more than the vanilla Gemini. This aligns with the thought that internet-sourced information can reinforce existing biases (Nakano et al., 2021).

In contrast to BLEnD, the non-selective KB grounding strategy significantly improved OLMo’s performance in this task because, when presented with several texts that seem irrelevant to the question, OLMo



could not ground any choice with the retrieved texts and therefore answered “unsure”, which was considered a correct answer for this task. By contrast, Gemini may be able to ignore irrelevant texts and exploit the relevant ones, resulting in stereotypical answers.

To understand how stereotypical facts in the KB affected KB grounding, we examined the sources of the retrieved texts and found that, out of 4,600 questions, in the KB query “without choices” setting, 1,156 questions retrieved at least one SeeGULL stereotype. However, in only 35 of these did the stereotype’s content exactly match the question. Similarly, in the KB query “with choices” setting, 1,266 questions retrieved at least one SeeGULL stereotype, but the stereotype’s content exactly matched the question in only 2 cases. In both settings, SeeGULL stereotypes whose content matched the question always passed the relevancy check for GPT and OLMo and usually passed for Gemini (~80% of the time). The inclusion of these stereotypes in the prompt flipped the original “unsure” answer of the vanilla model to a stereotypical answer. This suggests that including stereotypes in a prompt can induce a model to affirm stereotypes. However, as the KB has significantly fewer and narrower stereotypical contents than the internet, this issue is not as severe as the search-grounding strategy.

## 4 Human Evaluation

**Setup.** To evaluate the strategies on a more open-ended text generation task, we translated five questions from BLEND and five questions from NORMAD into open-ended prompts asking the model to tell a story set in a particular country. We adapted each of these ten questions for the ten national cultures used in the BLEND benchmark evaluation, leaving us with 100 (country, prompt) pairs. Next, we generated responses from Gemini using the vanilla, the selective KB-grounding, and the search-grounding. For each strategy, we generated three unique responses with the temperature of 0.5. Then we recruited nine evaluators from each of the ten studied cultures to rate, on a scale from 0 to 4, how culturally familiar each response was and to provide a brief justification of their score (see Appendices A.6–A.7 for details).

**Results.** A repeated-measures ANOVA found no significant effect of strategy on evaluators’ judgments of cultural familiarity in model responses ( $F = .18, p = .827$ ). Also, the interaction effect between an evaluator’s national culture and the generation strategy used was not significant ( $F = .84, p = .651$ ). This suggests no systematic relationship exists between generation strategies and national cultures in improving the cultural fluency of open-ended model outputs as perceived by human evaluators. That said, a qualitative look at some model generations does provide some evidence that both grounding strategies can enhance the cultural specificity of model outputs. In response to the prompt ‘Tell me a story in Mexico in which a

group of people of varying ages eat together and all guests behave in a socially acceptable way,’ the selective KB-grounded and search-grounded responses each mentioned many specific dishes and games, while the vanilla response was far more generic (see Appendix A.8 for examples). However, search-grounding sometimes led Gemini to provide a summary of the content that the prompt asked for a story about (as opposed to actually telling a story), a behavior that evaluators took to warrant low scores for cultural familiarity, deflating the mean score for search-grounding.

## 5 Discussion and Conclusion

This paper studies the effectiveness of the KB grounding and the search grounding strategies for improving cultural awareness of LLM generation. We discuss the key findings and suggestions below.

**KB Grounding vs Search Grounding.** The advantages of search-grounding on BLEND and NORMAD speak to the vast space of cultural facts on the internet. Even the large KB we compiled here still lacked many culturally relevant facts. We also observed some bias in each knowledge source used: approximately 19% of CultureAtlas entries and 25% of CultureBank entries concern the culture of the United States. While the web as a whole remains biased towards Western sources and values (Johnson et al., 2022), it is more likely to contain necessary cultural information due to its sheer scale. However, the poor performance of search-grounding on the stereotype avoidance benchmark reminds us that the context retrieved via web search could reinforce the (typically false) notion that stereotypes are factual and encourage models to affirm those stereotypes. This suggests that search-grounding is not yet a panacea for improving the cultural sensitivity of LLMs.

**Knowledge vs Fluency.** The results of the human evaluation do not show that search-grounding or KB-grounding improved the cultural familiarity of LLM outputs. The divergence in the performance of both strategies on the human evaluation and the multiple-choice QA benchmarks suggests that we ought to draw a distinction between two varieties of cultural awareness. On the one hand, there is a sense of cultural awareness that involves possessing **propositional knowledge** about a culture (i.e., knowing facts about that culture), as measured by the multiple-choice QA benchmarks, where the grounding strategies can improve the model performance. However, another sense of cultural awareness involves writing and speaking like someone with first-hand experience of and immersion in a culture, i.e., a sense of **cultural fluency**, as measured by our human evaluation, which revealed that the two grounding strategies were of limited value. We leave it to future work to develop strategies for improving the cultural awareness of generative language models along this second axis.

## Limitations

We acknowledge the several limitations of our work in this paper. First, we ran our evaluations only on smaller versions of the GPT-4, Gemini 1.5, and OLMO 2 models. It remains an open question whether our pattern of results would be similar for larger versions of these models and other model families.

Second, for BLEnD and the human evaluation, we ran our evaluations using prompts that were relevant to ten national countries and cultures; a more comprehensive study would require the use of a wider range of national *and* regional cultures. With the momentum in the community to create more culturally salient resources, a more comprehensive study in the future will help identify gaps and interventions for the majority world.

Third, we only implement the search-grounding strategy using the Gemini model, specifically using the “Grounding with Google Search” feature of Vertex AI. According to the metadata returned, this end-to-end API has its own methods for query rewriting, incorporating retrieved texts, and providing citations, the details of which are not publicly available. Since these steps are tied to Gemini, it is not replicable for other LLMs while maintaining a fair comparison. As more LLM APIs introduce search-grounding capabilities, we hope that the lessons learned from search-grounding Gemini, as presented in this paper, establish the necessity of carefully auditing the characteristics of any future search-grounded LLMs before using them in culturally sensitive applications. Moreover, the paper retains its core message that RAG (including search-grounding) can improve performance with respect to cultural propositional knowledge, but not necessarily cultural fluency.

Finally, all of our evaluations concerned solely English-language prompts and outputs. While we take it to be an important goal for generative language models that they be able to produce culturally-aware outputs about any culture in the world’s most widely-spoken language, the landscape of cultural awareness becomes much more nuanced when one considers the rich variegation that exists in phrasing and dialect across a wide range of languages. We leave it to future work to examine whether the strategies used here can be adapted to a multi-lingual context.

## Ethical Considerations

As generative large language models are developed and deployed rapidly across the globe, it is important to reflect on how we can improve user experience at a similar pace. The promise of model utility for a myriad of tasks such as that of a writing assistant, remains unfulfilled if the model is not beneficial or usable for a vast majority. With this work, we attempted to begin adaptation of techniques in NLP to further the cause of cultural awareness and relevance in models. As noted in our limitations, with more comprehensive work across

a greater number of cultures and countries, we hope that development of more culturally-aware models will be possible.

## References

- Muhammad Farid Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Shivdutt Singh, Alham Fikri Aji, Jacki O’Neill, Ashutosh Modi, and Monojit Choudhury. 2024. [Towards measuring and modeling “culture” in LLMs: A survey](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 15763–15784, Miami, Florida, USA. Association for Computational Linguistics.
- Dhruv Agarwal, Mor Naaman, and Aditya Vashistha. 2025. [Ai suggestions homogenize writing toward western styles and diminish cultural nuances](#).
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). *arXiv preprint arXiv:2310.11511*.
- Mohammad Atari, Mona J Xue, Peter S Park, Damián Blasi, and Joseph Henrich. 2023. [Which humans?](#) *PsyArXiv*.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *arXiv preprint arXiv:2204.05862*.
- Mukul Bhutani, Kevin Robinson, Vinodkumar Prabhakaran, Shachi Dave, and Sunipa Dev. 2024. [SeeGULL multilingual: a dataset of geo-culturally situated stereotypes](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 842–854, Bangkok, Thailand. Association for Computational Linguistics.
- Alex J Chan, José Luis Redondo García, Fabrizio Silvestri, Colm O’Donnell, and Konstantina Palla. 2023. [Enhancing content moderation with culturally-aware models](#). *arXiv e-prints*, pages arXiv–2312.
- Jesse Dodge, Maarten Sap, Ana Marasović, William Agnew, Gabriel Ilharco, Dirk Groeneveld, Margaret Mitchell, and Matt Gardner. 2021. [Documenting large webtext corpora: A case study on the colossal clean crawled corpus](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1286–1305, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Wenqi Fan, Yujuan Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and Qing Li. 2024. [A survey on rag meeting llms: Towards retrieval-augmented large language models](#).

- In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 6491–6501.
- Yi Fung, Ruining Zhao, Jae Doo, Chenkai Sun, and Heng Ji. 2024. Massively multi-cultural knowledge acquisition & lm benchmarking. *arXiv preprint arXiv:2402.09369*.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. [Retrieval-augmented generation for large language models: A survey](#). *arXiv preprint arXiv:2312.10997*.
- Akshita Jha, Aida Mostafazadeh Davani, Chandan K Reddy, Shachi Dave, Vinodkumar Prabhakaran, and Sunipa Dev. 2023. [SeeGULL: A stereotype benchmark with broad geo-cultural coverage leveraging generative models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9851–9870, Toronto, Canada. Association for Computational Linguistics.
- Rebecca L Johnson, Giada Pistilli, Natalia Menéndez-González, Leslye Denisse Dias Duran, Enrico Panai, Julija Kalpokiene, and Donald Jay Bertulfo. 2022. The ghost in the machine has an american accent: value conflict in gpt-3. *arXiv preprint arXiv:2203.07785*.
- Hailey Joren, Jianyi Zhang, Chun-Sung Ferng, Da-Cheng Juan, Ankur Taly, and Cyrus Rashtchian. 2024. Sufficient context: A new lens on retrieval augmented generation systems. *arXiv preprint arXiv:2411.06037*.
- Nithish Kannan, Arif Ahmad, marco Andreetto, Vinodkumar Prabhakaran, Utsav Prabhu, Adji Bousso Dieng, Pushpak Bhattacharyya, and Shachi Dave. 2024. [Beyond aesthetics: Cultural competence in text-to-image models](#). In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Mohammed Abdul Khaliq, Paul Yu-Chun Chang, Mingyang Ma, Bernhard Pflugfelder, and Filip Miletic. 2024. [RAGAR, your falsehood radar: RAG-augmented reasoning for political fact-checking using multimodal large language models](#). In *Proceedings of the Seventh Fact Extraction and VERification Workshop (FEVER)*, pages 280–296, Miami, Florida, USA. Association for Computational Linguistics.
- Jaewoong Kim and Moohong Min. 2024. From rag to qa-rag: Integrating generative ai for pharmaceutical regulatory compliance process. *arXiv preprint arXiv:2402.01717*.
- Mojtaba Komeili, Kurt Shuster, and Jason Weston. 2022. [Internet-augmented dialogue generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8460–8478, Dublin, Ireland. Association for Computational Linguistics.
- Angeliki Lazaridou, Elena Gribovskaya, Wojciech Stokowiec, and Nikolai Grigorev. 2022. Internet-augmented language models through few-shot prompting for open-domain question answering. *arXiv preprint arXiv:2203.05115*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Cheng Li, Mengzhou Chen, Jindong Wang, Sunayana Sitaram, and Xing Xie. 2024a. Culturellm: Incorporating cultural differences into large language models. In *Thirty-Eighth Annual Conference on Neural Information Processing Systems (NeurIPS)*.
- Cheng Li, Damien Teney, Linyi Yang, Qingsong Wen, Xing Xie, and Jindong Wang. 2024b. Culturepark: Boosting cross-cultural understanding in large language models. In *Thirty-Eighth Annual Conference on Neural Information Processing Systems (NeurIPS)*.
- Jiarui Li, Ye Yuan, and Zehua Zhang. 2024c. Enhancing llm factual accuracy with rag to counter hallucinations: A case study on domain-specific queries in private knowledge-bases. *arXiv preprint arXiv:2403.10446*.
- Junho Myung, Nayeon Lee, Yi Zhou, Jiho Jin, Rifki Afina Putri, Dimosthenis Antypas, Hsuvas Borkakoty, Eunsu Kim, Carla Perez-Almendros, Abinew Ali Ayele, Víctor Gutiérrez-Basulto, Yazmín Ibáñez-García, Hwaran Lee, Sham-suddeen Hassan Muhammad, Kiwoong Park, Anar Sabuhi Rzayev, Nina White, Seid Muhie Yimam, Mohammad Taher Pilehvar, Nedjma Ousidhoum, Jose Camacho-Collados, and Alice Oh. 2024. [Blend: A benchmark for llms on everyday knowledge in diverse cultures and languages](#).
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*.
- Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafjord, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, Michal Guerquin, Hamish Ivison, Pang Wei Koh, Jiacheng Liu, Saumya Malik, William Merrill, Lester James V. Miranda, Jacob Morrison, Tyler Murray, Crystal Nam, Valentina Pyatkin, Aman Rangapur, Michael Schmitz, Sam Skjonsberg, David Wadden, Christopher Wilhelm, Michael Wilson, Luke Zettlemoyer, Ali Farhadi, Noah A. Smith, and Hannaneh Hajishirzi. 2025. [2 olmo 2 furious](#).

OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mađry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Conneau, Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian, Amin Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein, Andrew Cann, Andrew Codispoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrey Mishchenko, Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia, Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben Leimberger, Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake Samic, Bob McGrew, Bobby Spero, Bogo Giertler, Bowen Cheng, Brad Lightcap, Brandon Walkin, Brendan Quinn, Brian Guarraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo Lugaresi, Carroll Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li, Chan Jun Shern, Channing Conger, Charlotte Barette, Chelsea Voss, Chen Ding, Cheng Lu, Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim, Christine Choi, Christine McLeavey, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, Dane Sherburn, Daniel Kappler, Daniel Levin, Daniel Levy, David Carr, David Farhi, David Mely, David Robinson, David Sasaki, Denny Jin, Dev Valadares, Dimitris Tsipras, Doug Li, Duc Phong Nguyen, Duncan Findlay, Edede Oiwoh, Edmund Wong, Ehsan Asdar, Elizabeth Proehl, Elizabeth Yang, Eric Antonow, Eric Kramer, Eric Peterson, Eric Sigler, Eric Wallace, Eugene Brevdo, Evan Mays, Farzad Khorasani, Felipe Petroski Such, Filippo Raso, Francis Zhang, Fred von Lohmann, Freddie Sulit, Gabriel Goh, Gene Oden, Geoff Salmon, Giulio Starace, Greg Brockman, Hadi Salman, Haiming Bao, Haitang Hu, Hannah Wong, Haoyu Wang, Heather Schmidt, Heather Whitney, Heewoo Jun, Hendrik Kirchner, Henrique Ponde de Oliveira Pinto, Hongyu Ren, Huiwen Chang, Hyung Won Chung, Ian Kivlichan, Ian O'Connell, Ian O'Connell, Ian Osband, Ian Silber, Ian Sohl, Ibrahim Okuyucu, Ikai Lan, Ilya Kostrikov, Ilya Sutskever, Ingmar Kanitscheider, Ishaan Gulrajani, Jacob Coxon, Jacob Menick, Jakub Pachocki, James Aung, James Betker, James Crooks, James Lennon, Jamie Kiros, Jan Leike, Jane Park, Jason Kwon, Jason Phang, Jason Teplitz, Jason Wei, Jason Wolfe, Jay Chen, Jeff Harris, Jenia Varavva, Jessica Gan Lee, Jessica Shieh, Ji Lin, Jiahui Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joanne Jang, Joaquin Quinonero Candela, Joe Beutler, Joe Landers, Joel Parish, Johannes Heidecke, John Schulman, Jonathan Lachman, Jonathan McKay, Jonathan Uesato, Jonathan Ward, Jong Wook Kim, Joost Huizinga, Jordan Sitkin, Jos Kraaijeveld, Josh

Gross, Josh Kaplan, Josh Snyder, Joshua Achiam, Joy Jiao, Joyce Lee, Juntang Zhuang, Justyn Harriman, Kai Fricke, Kai Hayashi, Karan Singhal, Katy Shi, Kavin Karthik, Kayla Wood, Kendra Rimbach, Kenny Hsu, Kenny Nguyen, Keren Gulemberg, Kevin Button, Kevin Liu, Kiel Howe, Krithika Muthukumar, Kyle Luther, Lama Ahmad, Larry Kai, Lauren Itow, Lauren Workman, Leher Pathak, Leo Chen, Li Jing, Lia Guy, Liam Fedus, Liang Zhou, Lien Mamitsuka, Lilian Weng, Lindsay McCallum, Lindsey Held, Long Ouyang, Louis Feувrier, Lu Zhang, Lukas Kondraciuk, Lukasz Kaiser, Luke Hewitt, Luke Metz, Lyric Doshi, Mada Aflak, Maddie Simens, Madelaine Boyd, Madeleine Thompson, Marat Dukhan, Mark Chen, Mark Gray, Mark Hudnall, Marvin Zhang, Marwan Aljube, Mateusz Litwin, Matthew Zeng, Max Johnson, Maya Shetty, Mayank Gupta, Meghan Shah, Mehmet Yatbaz, Meng Jia Yang, Mengchao Zhong, Mia Glaese, Mianna Chen, Michael Janner, Michael Lampe, Michael Petrov, Michael Wu, Michele Wang, Michelle Fradin, Michelle Pokrass, Miguel Castro, Miguel Oom Temudo de Castro, Mikhail Pavlov, Miles Brundage, Miles Wang, Minal Khan, Mira Murati, Mo Bavarian, Molly Lin, Murat Yesildal, Nacho Soto, Natalia Gimelshein, Natalie Cone, Natalie Staudacher, Natalie Summers, Natan LaFontaine, Neil Chowdhury, Nick Ryder, Nick Stathas, Nick Turley, Nik Tezak, Niko Felix, Nithanth Kudige, Nitish Keskar, Noah Deutsch, Noel Bundick, Nora Puckett, Ofir Nachum, Ola Okelola, Oleg Boiko, Oleg Murk, Oliver Jaffe, Olivia Watkins, Olivier Godement, Owen Campbell-Moore, Patrick Chao, Paul McMillan, Pavel Belov, Peng Su, Peter Bak, Peter Bakkum, Peter Deng, Peter Dolan, Peter Hoeschele, Peter Welinder, Phil Tillet, Philip Pronin, Philippe Tillet, Prafulla Dhariwal, Qiming Yuan, Rachel Dias, Rachel Lim, Rahul Arora, Rajan Troll, Randall Lin, Rapha Gontijo Lopes, Raul Puri, Reah Miyara, Reimar Leike, Renaud Gaubert, Reza Zamani, Ricky Wang, Rob Donnelly, Rob Honsby, Rocky Smith, Rohan Sahai, Rohit Ramchandani, Romain Huet, Rory Carmichael, Rowan Zellers, Roy Chen, Ruby Chen, Ruslan Nigmatullin, Ryan Cheu, Saachi Jain, Sam Altman, Sam Schoenholz, Sam Toizer, Samuel Miserendino, Sandhini Agarwal, Sara Culver, Scott Ethersmith, Scott Gray, Sean Grove, Sean Metzger, Shamez Hermani, Shantanu Jain, Shengjia Zhao, Sherwin Wu, Shino Jomoto, Shirong Wu, Shuaiqi, Xia, Sonia Phene, Spencer Papay, Srinivas Narayanan, Steve Coffey, Steve Lee, Stewart Hall, Suchir Balaaji, Tal Broda, Tal Stramer, Tao Xu, Tarun Gogineni, Taya Christianson, Ted Sanders, Tejal Patwardhan, Thomas Cunningham, Thomas Degry, Thomas Dimson, Thomas Raoux, Thomas Shadwell, Tianhao Zheng, Todd Underwood, Todor Markov, Toki Sherbakov, Tom Rubin, Tom Stasi, Tomer Kaftan, Tristan Heywood, Troy Peterson, Tyce Walters, Tyna Eloundou, Valerie Qi, Veit Moeller, Vinnie Monaco, Vishal Kuo, Vlad Fomenko, Wayne Chang, Weiyi Zheng, Wenda Zhou, Wesam Manassra, Will Sheu, Wojciech Zaremba, Yash Patil, Yilei Qian, Yongjik Kim, Youlong Cheng, Yu Zhang, Yuchen



- He, Yuchen Zhang, Yujia Jin, Yunxing Dai, and Yury Malkov. 2024. [Gpt-4o system card](#).
- Siddhesh Pawar, Junyeong Park, Jiho Jin, Arnav Arora, Junho Myung, Srishti Yadav, Faiz Ghifari Haznitrama, Inhwa Song, Alice Oh, and Isabelle Augenstein. 2024. [Survey of cultural awareness in language models: Text and beyond](#).
- Vinodkumar Prabhakaran, Rida Qadri, and Ben Hutchinson. 2022. Cultural incongruencies in artificial intelligence. *arXiv preprint arXiv:2211.13069*.
- Rida Qadri, Aida M. Davani, Kevin Robinson, and Vinodkumar Prabhakaran. 2025. [Risks of cultural erasure in large language models](#).
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Abhinav Rao, Akhila Yerukola, Vishwa Shah, Katharina Reinecke, and Maarten Sap. 2024. [Normad: A framework for measuring the cultural adaptability of large language models](#).
- Abhinav Sukumar Rao, Aditi Khandelwal, Kumar Tanmay, Utkarsh Agarwal, and Monojit Choudhury. 2023. [Ethical reasoning over moral alignment: A case and framework for in-context ethical policies in LLMs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13370–13388, Singapore. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Minju Seo, Jinheon Baek, James Thorne, and Sung Ju Hwang. 2024. Retrieval-augmented data augmentation for low-resource domain tasks. *arXiv preprint arXiv:2402.13482*.
- Weiyang Shi, Ryan Li, Yutong Zhang, Caleb Ziems, Raya Horesh, Rogério Abreu de Paula, Diyi Yang, et al. 2024. [Culturebank: An online community-driven knowledge base towards culturally aware language technologies](#). *arXiv preprint arXiv:2404.15238*.
- Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, et al. 2022. [Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage](#). *arXiv preprint arXiv:2208.03188*.
- Ronit Singal, Pransh Patwa, Parth Patwa, Aman Chadha, and Amitava Das. 2024. [Evidence-backed fact checking using RAG and few-shot in-context learning with LLMs](#). In *Proceedings of the Seventh Fact Extraction and VERification Workshop (FEVER)*, pages 91–98, Miami, Florida, USA. Association for Computational Linguistics.
- Binhuan Sun, Liubov Pashkova, Pascal Aldo Pieters, Archana Sanjay Harke, Omkar Satyavan Mohite, Alberto Santos, Daniel C Zielinski, Bernhard O Pals-son, and Patrick Victor Phaneuf. 2025. [Pankb: An interactive microbial pangenome knowledge-base for research, biotechnological innovation, and knowledge mining](#). *Nucleic Acids Research*, 53(D1):D806–D818.
- Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, Soroosh Mariooryad, Yifan Ding, Xinyang Geng, Fred Alcober, Roy Frostig, Mark Omernick, Lexi Walker, Cosmin Paduraru, Christina Sorokin, Andrea Tacchetti, Colin Gaffney, Samira Daruki, Olcan Sercinoglu, Zach Gleicher, Juliette Love, Paul Voigtlaender, Rohan Jain, Gabriela Surita, Kareem Mohamed, Rory Blevins, Junwhan Ahn, Tao Zhu, Kornraphop Kawintiranon, Orhan Firat, Yiming Gu, Yujing Zhang, Matthew Rahtz, Manaal Faruqui, Natalie Clay, Justin Gilmer, JD Co-Reyes, Ivo Penchev, Rui Zhu, Nobuyuki Morioka, Kevin Hui, Krishna Haridasan, Victor Campos, Mahdis Mahdieh, Mandy Guo, Samer Hassan, Kevin Kilgour, Arpi Vezer, Heng-Tze Cheng, Raoul de Liedekerke, Siddharth Goyal, Paul Barham, DJ Strouse, Seb Noury, Jonas Adler, Mukund Sundararajan, Sharad Vikram, Dmitry Lepikhin, Michela Paganini, Xavier Garcia, Fan Yang, Dasha Valter, Maja Trebacz, Kiran Vodrahalli, Chulayuth Asawaroengchai, Roman Ring, Norbert Kalb, Livio Baldini Soares, Siddhartha Brahma, David Steiner, Tianhe Yu, Fabian Mentzer, Antoine He, Lucas Gonzalez, Bibo Xu, Raphael Lopez Kaufman, Laurent El Shafey, Junhyuk Oh, Tom Hennigan, George van den Driessche, Seth Odoom, Mario Lucic, Becca Roelofs, Sid Lall, Amit Marathe, Betty Chan, Santiago Ontanon, Luheng He, Denis Teplyashin, Jonathan Lai, Phil Crone, Bogdan Damoc, Lewis Ho, Sebastian Riedel, Karel Lenc, Chih-Kuan Yeh, Aakanksha Chowdhery, Yang Xu, Mehran Kazemi, Ehsan Amid, Anastasia Petrushkina, Kevin Swersky, Ali Khodaei, Gwoon Chen, Chris Larkin, Mario Pinto, Geng Yan, Adria Puigdomenech Badia, Piyush Patil, Steven Hansen, Dave Orr, Sebastien M. R. Arnold, Jordan Grimstad, Andrew Dai, Sholto Douglas, Rishika Sinha, Vikas Yadav, Xi Chen, Elena Gribovskaya, Jacob Austin, Jeffrey Zhao, Kaushal Patel, Paul Komarek, Sophia Austin, Sebastian Borgeaud, Linda Friso, Abhimanyu Goyal, Ben Caine, Kris Cao, Da-Woon Chung, Matthew Lamm, Gabe Barth-Maron, Thais Kagohara, Kate Olaszewska, Mia Chen, Kaushik Shivakumar, Rishabh Agarwal, Harshal Godhia, Ravi Rajwar, Javier Snaider, Xerxes Dotiwalla, Yuan Liu, Aditya Barua, Victor Ungureanu, Yuan Zhang, Bat-Orgil Batsaikhan, Mateo Wirth, James Qin, Ivo Danihelka, Tulsee Doshi, Martin Chadwick, Jilin Chen, Sanil Jain, Quoc Le, Arjun Kar, Madhu Gurusurthy,

Cheng Li, Ruoxin Sang, Fangyu Liu, Lampros Lamprou, Rich Munoz, Nathan Lintz, Harsh Mehta, Heidi Howard, Malcolm Reynolds, Lora Aroyo, Quan Wang, Lorenzo Blanco, Albin Cassirer, Jordan Griffith, Dipanjan Das, Stephan Lee, Jakub Szygnowski, Zach Fisher, James Besley, Richard Powell, Zafarali Ahmed, Dominik Paulus, David Ritter, Zalan Borsos, Rishabh Joshi, Aedan Pope, Steven Hand, Vittorio Selo, Vihan Jain, Nikhil Sethi, Megha Goel, Takaki Makino, Rhys May, Zhen Yang, Johan Schalkwyk, Christina Butterfield, Anja Hauth, Alex Goldin, Will Hawkins, Evan Senter, Sergey Brin, Oliver Woodman, Marvin Ritter, Eric Noland, Minh Giang, Vijay Bolina, Lisa Lee, Tim Blyth, Ian Mackinnon, Machel Reid, Obaid Sarvana, David Silver, Alexander Chen, Lily Wang, Loren Maggiore, Oscar Chang, Nithya Attaluri, Gregory Thornton, Chung-Cheng Chiu, Oskar Bunyan, Nir Levine, Timothy Chung, Evgenii Eltyshev, Xiance Si, Timothy Lillicrap, Demetra Brady, Vaibhav Aggarwal, Boxi Wu, Yuanzhong Xu, Ross McIlroy, Kartikeya Badola, Paramjit Sandhu, Erica Moreira, Wojciech Stokowiec, Ross Hemsley, Dong Li, Alex Tudor, Pranav Shyam, Elahe Rahimtoroghi, Salem Haykal, Pablo Sprechmann, Xiang Zhou, Diana Mincu, Yujia Li, Ravi Addanki, Kalpesh Krishna, Xiao Wu, Alexandre Frechette, Matan Eyal, Allan Dafoe, Dave Lacey, Jay Whang, Thi Avrahami, Ye Zhang, Emanuel Taropa, Hanzhao Lin, Daniel Toyama, Eliza Rutherford, Motoki Sano, HyunJeong Choe, Alex Tomala, Chalence Safranek-Shrader, Nora Kassner, Mantas Pajarskas, Matt Harvey, Sean Sechrist, Meire Fortunato, Christina Lyu, Gamaleldin Elsayed, Chenkai Kuang, James Lottes, Eric Chu, Chao Jia, Chih-Wei Chen, Peter Humphreys, Kate Baumli, Connie Tao, Rajkumar Samuel, Cicero Nogueira dos Santos, Anders Andreassen, Nemanja Rakićević, Dominik Grewe, Aviral Kumar, Stephanie Winkler, Jonathan Caton, Andrew Brock, Sid Dalmia, Hannah Sheahan, Iain Barr, Yingjie Miao, Paul Natsev, Jacob Devlin, Feryal Behbahani, Flavian Prost, Yanhua Sun, Artiom Myaskovsky, Thanumalayan Sankaranarayanan Pillai, Dan Hurt, Angeliki Lazaridou, Xi Xiong, Ce Zheng, Fabio Pardo, Xiaowei Li, Dan Horgan, Joe Stanton, Moran Ambar, Fei Xia, Alejandro Lince, Mingqiu Wang, Basil Mustafa, Albert Webson, Hyo Lee, Rohan Anil, Martin Wicke, Timothy Dozat, Abhishek Sinha, Enrique Piqueras, Elahe Dabir, Shyam Upadhyay, Anudhyan Borral, Lisa Anne Hendricks, Corey Fry, Josip Djolonga, Yi Su, Jake Walker, Jane Labanowski, Ronny Huang, Vedant Misra, Jeremy Chen, RJ Skerry-Ryan, Avi Singh, Shruti Rijhwani, Dian Yu, Alex Castro-Ros, Beer Changpinyo, Romina Datta, Sumit Bagri, Arnar Mar Hrafnkelsson, Marcello Maggioni, Daniel Zheng, Yury Sulsky, Shaobo Hou, Tom Le Paine, Antoine Yang, Jason Riesa, Dominika Rogozinska, Dror Marcus, Dalia El Badawy, Qiao Zhang, Luyu Wang, Helen Miller, Jeremy Greer, Lars Lowe Sjos, Azade Nova, Heiga Zen, Rahma Chaabouni, Mihaela Rosca, Jiepu Jiang, Charlie Chen, Ruibo Liu, Tara Sainath, Maxim Krikun, Alex Polozov, Jean-Baptiste Lespiau, Josh Newlan, Zeyn-

cep Cankara, Soo Kwak, Yunhan Xu, Phil Chen, Andy Coenen, Clemens Meyer, Katerina Tsihlias, Ada Ma, Juraj Gottweis, Jinwei Xing, Chenjie Gu, Jin Miao, Christian Frank, Zeynep Cankara, Sanjay Ganapathy, Ishita Dasgupta, Steph Hughes-Fitt, Heng Chen, David Reid, Keran Rong, Hongmin Fan, Joost van Amersfoort, Vincent Zhuang, Aaron Cohen, Shixiang Shane Gu, Anhad Mohananey, Anastasija Ilic, Taylor Tobin, John Wieting, Anna Bortsova, Phoebe Thacker, Emma Wang, Emily Caveness, Justin Chiu, Eren Sezener, Alex Kaskasoli, Steven Baker, Katie Millican, Mohamed Elhawaty, Kostas Aisopos, Carl Lebsack, Nathan Byrd, Hanjun Dai, Wenhao Jia, Matthew Wiethoff, Elnaz Davoodi, Albert Weston, Lakshman Yagati, Arun Ahuja, Isabel Gao, Golan Pundak, Susan Zhang, Michael Azzam, Khe Chai Sim, Sergi Caelles, James Keeling, Abhanshu Sharma, Andy Swing, YaGuang Li, Chenxi Liu, Carrie Grimes Bostock, Yamini Bansal, Zachary Nado, Anakesh Anand, Josh Lipschultz, Abhijit Karmarkar, Lev Proleev, Abe Ittycheriah, Soheil Hassas Yeganeh, George Polovets, Aleksandra Faust, Jiao Sun, Alban Rrustemi, Pen Li, Rakesh Shivanna, Jeremiah Liu, Chris Welty, Federico Lebron, Anirudh Baddepudi, Sebastian Krause, Emilio Parisotto, Radu Soricut, Zheng Xu, Dawn Bloxwich, Melvin Johnson, Behnam Neyshabur, Justin Mao-Jones, Renshen Wang, Vinay Ramasesh, Zaheer Abbas, Arthur Guez, Constant Segal, Duc Dung Nguyen, James Svensson, Le Hou, Sarah York, Kieran Milan, Sophie Bridgers, Wiktor Gworek, Marco Tagliasacchi, James Lee-Thorp, Michael Chang, Alexey Guseynov, Ale Jakse Hartman, Michael Kwong, Ruizhe Zhao, Sheleem Kashem, Elizabeth Cole, Antoine Miech, Richard Tanburn, Mary Phuong, Filip Pavetic, Sebastien Cevey, Ramona Comanescu, Richard Ives, Sherry Yang, Cosmo Du, Bo Li, Zizhao Zhang, Mariko Iinuma, Clara Huiyi Hu, Aurko Roy, Shaan Bijwadia, Zhenkai Zhu, Danilo Martins, Rachel Saputro, Anita Gergely, Steven Zheng, Dawei Jia, Ioannis Antonoglou, Adam Sadovsky, Shane Gu, Yingying Bi, Alek Andreev, Sina Samangooei, Mina Khan, Tomas Kocisky, Angelos Filos, Chintu Kumar, Colton Bishop, Adams Yu, Sarah Hodkinson, Sid Mittal, Premal Shah, Alexandre Moufarek, Yong Cheng, Adam Bloniarz, Jaehoon Lee, Pedram Pejman, Paul Michel, Stephen Spencer, Vladimir Feinberg, Xuehan Xiong, Nikolay Savinov, Charlotte Smith, Siamak Shakeri, Dustin Tran, Mary Chesus, Bernd Bohnet, George Tucker, Tamara von Glehn, Carrie Muir, Yiran Mao, Hideto Kazawa, Ambrose Slone, Kedar Soparkar, Disha Shrivastava, James Cobon-Kerr, Michael Sharman, Jay Pavagadhi, Carlos Araya, Karolis Misiunas, Nimesh Ghelani, Michael Laskin, David Barker, Qijia Li, Anton Briukhov, Neil Houlsby, Mia Glaese, Balaji Lakshminarayanan, Nathan Schucher, Yunhao Tang, Eli Collins, Hyeontaek Lim, Fangxiaoyu Feng, Adria Recasens, Guangda Lai, Alberto Magni, Nicola De Cao, Aditya Siddhant, Zoe Ashwood, Jordi Orbay, Mostafa Dehghani, Jenny Brennan, Yifan He, Kelvin Xu, Yang Gao, Carl Saroufim, James Molloy, Xinyi Wu, Seb Arnold, Solomon

Chang, Julian Schrittwieser, Elena Buchatskaya, Soroush Radpour, Martin Polacek, Skye Giordano, Ankur Bapna, Simon Tokumine, Vincent Helledoorn, Thibault Sottiaux, Sarah Cogan, Aliaksei Severyn, Mohammad Saleh, Shantanu Thakoor, Laurent Shefey, Siyuan Qiao, Meenu Gaba, Shuo yiin Chang, Craig Swanson, Biao Zhang, Benjamin Lee, Paul Kishan Rubenstein, Gan Song, Tom Kwiatkowski, Anna Koop, Ajay Kannan, David Kao, Parker Schuh, Axel Stjerngren, Golnaz Ghiasi, Gena Gibson, Luke Vilnis, Ye Yuan, Felipe Tiengo Ferreira, Aishwarya Kamath, Ted Klimentko, Ken Franko, Kefan Xiao, Indro Bhattacharya, Miteyan Patel, Rui Wang, Alex Morris, Robin Strudel, Vivek Sharma, Peter Choy, Sayed Hadi Hashemi, Jessica Landon, Mara Finkelstein, Priya Jhakra, Justin Frye, Megan Barnes, Matthew Mauger, Dennis Daun, Khuslen Baatarsukh, Matthew Tung, Wael Farhan, Henryk Michalewski, Fabio Viola, Felix de Chaumont Quiry, Charline Le Lan, Tom Hudson, Qingze Wang, Felix Fischer, Ivy Zheng, Elspeth White, Anca Dragan, Jean baptiste Alayrac, Eric Ni, Alexander Pritzel, Adam Iwanicki, Michael Isard, Anna Bulanova, Lukas Zilka, Ethan Dyer, Devendra Sachan, Srivatsan Srinivasan, Hannah Muckenhirn, Honglong Cai, Amol Mandhane, Mukarram Tariq, Jack W. Rae, Gary Wang, Kareem Ayoub, Nicholas FitzGerald, Yao Zhao, Woo Hyun Han, Chris Alberti, Dan Garrette, Kashyap Krishnakumar, Mai Gimenez, Anselm Levskaia, Daniel Sohn, Josip Matak, Inaki Iturrate, Michael B. Chang, Jackie Xiang, Yuan Cao, Nishant Ranka, Geoff Brown, Adrian Hutter, Vahab Mirrokni, Nanxin Chen, Kaisheng Yao, Zoltan Egyed, Francois Galilee, Tyler Liechty, Praveen Kallakuri, Evan Palmer, Sanjay Ghemawat, Jasmine Liu, David Tao, Chloe Thornton, Tim Green, Mimi Jasarevic, Sharon Lin, Victor Cotruta, Yi-Xuan Tan, Noah Fiedel, Hongkun Yu, Ed Chi, Alexander Neitz, Jens Heitkaemper, Anu Sinha, Denny Zhou, Yi Sun, Charbel Kaed, Brice Hulse, Swaroop Mishra, Maria Georgaki, Sneha Kudugunta, Clement Farabet, Izhak Shafran, Daniel Vlasic, Anton Tsitsulin, Rajagopal Ananthanarayanan, Alen Carin, Guolong Su, Pei Sun, Shashank V, Gabriel Carvajal, Josef Broder, Iulia Comsa, Alena Repina, William Wong, Warren Weilun Chen, Peter Hawkins, Egor Filonov, Lucia Loher, Christoph Hirschall, Weiwei Wang, Jingchen Ye, Andrea Burns, Hardie Cate, Diana Gage Wright, Federico Piccinini, Lei Zhang, Chu-Cheng Lin, Ionel Gog, Yana Kulizhskaya, Ashwin Sreevatsa, Shuang Song, Luis C. Cobo, Anand Iyer, Chetan Tekur, Guillermo Garrido, Zhuyun Xiao, Rupert Kemp, Huaixiu Steven Zheng, Hui Li, Ananth Agarwal, Christel Ngani, Kati Goshvadi, Rebeca Santamaria-Fernandez, Wojciech Fica, Xinyun Chen, Chris Gorgolewski, Sean Sun, Roopal Garg, Xinyu Ye, S. M. Ali Eslami, Nan Hua, Jon Simon, Pratik Joshi, Yelin Kim, Ian Tenney, Sahitya Potluri, Lam Nguyen Thiet, Quan Yuan, Florian Luisier, Alexandra Chronopoulou, Salvatore Scellato, Praveen Srinivasan, Minmin Chen, Vinod Koverkathu, Valentin Dalibard, Yaming Xu, Brennan Saeta, Keith Anderson, Thibault

Sellam, Nick Fernando, Fantine Huot, Junehyuk Jung, Mani Varadarajan, Michael Quinn, Amit Raul, Maigo Le, Ruslan Habalov, Jon Clark, Komal Jalan, Karesha Bullard, Achintya Singhal, Thang Luong, Boyu Wang, Sujevan Rajayogam, Julian Eisenschlos, Johnson Jia, Daniel Finchelstein, Alex Yakubovich, Daniel Balle, Michael Fink, Sameer Agarwal, Jing Li, Dj Dvijotham, Shalini Pal, Kai Kang, Jaclyn Konzelmann, Jennifer Beattie, Olivier Dousse, Diane Wu, Remi Crocker, Chen Elkind, Siddhartha Reddy Jonnalagadda, Jong Lee, Dan Holtmann-Rice, Krystal Kallarackal, Rosanne Liu, Denis Vnukov, Neera Vats, Luca Invernizzi, Mohsen Jafari, Huanjie Zhou, Lilly Taylor, Jennifer Prendki, Marcus Wu, Tom Eccles, Tianqi Liu, Kavya Koppapu, Francoise Beaufays, Christof Angermueller, Andreea Marzoca, Shourya Sarcar, Hilal Dib, Jeff Stanway, Frank Perbet, Nejc Trdin, Rachel Sterneck, Andrey Khorlin, Dinghua Li, Xihui Wu, Sonam Goenka, David Madras, Sasha Goldshtein, Willi Gierke, Tong Zhou, Yaxin Liu, Yannie Liang, Anais White, Yunjie Li, Shreya Singh, Sanaz Bahargam, Mark Epstein, Sujoy Basu, Li Lao, Adnan Ozturel, Carl Crous, Alex Zhai, Han Lu, Zora Tung, Neeraj Gaur, Alanna Walton, Lucas Dixon, Ming Zhang, Amir Globerson, Grant Uy, Andrew Bolt, Olivia Wiles, Milad Nasr, Ilia Shumailov, Marco Selvi, Francesco Piccinno, Ricardo Aguilar, Sara McCarthy, Misha Khalman, Mrinal Shukla, Vlado Galic, John Carpenter, Kevin Vellela, Haibin Zhang, Harry Richardson, James Martens, Matko Bosnjak, Shreyas Rammohan Belle, Jeff Seibert, Mahmoud Alnahlawi, Brian McWilliams, Sankalp Singh, Annie Louis, Wen Ding, Dan Popovici, Lenin Simicich, Laura Knight, Pulkit Mehta, Nishesh Gupta, Chongyang Shi, Saaber Fatehi, Jovana Mitrovic, Alex Grills, Joseph Pagadora, Tsendsuren Munkhdalai, Dessie Petrova, Danielle Eisenbud, Zhishuai Zhang, Damion Yates, Bhavishya Mittal, Nilesh Tripuraneni, Yannis Assael, Thomas Brovelli, Prateek Jain, Mihajlo Velimirovic, Canfer Akbulut, Jiaqi Mu, Wolfgang Macherey, Ravin Kumar, Jun Xu, Haroon Qureshi, Gheorghe Comanici, Jeremy Wiesner, Zhitao Gong, Anton Ruddock, Matthias Bauer, Nick Felt, Anirudh GP, Anurag Arnab, Dustin Zelle, Jonas Rothfuss, Bill Rosgen, Ashish Shenoy, Bryan Seybold, Xinjian Li, Jayaram Mudigonda, Goker Erdogan, Jiawei Xia, Jiri Simsa, Andrea Michi, Yi Yao, Christopher Yew, Steven Kan, Isaac Caswell, Carey Radebaugh, Andre Elisseeff, Pedro Valenzuela, Kay McKinney, Kim Paterson, Albert Cui, Eri Latorre-Chimoto, Solomon Kim, William Zeng, Ken Durden, Priya Ponnappalli, Tiberiu Sosea, Christopher A. Choquette-Choo, James Manyika, Brona Robenek, Harsha Vashisht, Sebastien Pereira, Hoi Lam, Marko Velic, Denese Owusu-Afriyie, Katherine Lee, Tolga Bolukbasi, Alicia Parrish, Shawn Lu, Jane Park, Balaji Venkatraman, Alice Talbert, Lambert Rosique, Yuchung Cheng, Andrei Sozanschi, Adam Paszke, Praveen Kumar, Jessica Austin, Lu Li, Khalid Salama, Bartek Perz, Wooyeol Kim, Nandita Dukupati, Anthony Baryshnikov, Christos Kaplanis, XiangHai Sheng, Yuri Chervonyi, Caglar Unlu, Diego de Las Casas, Harry

Askham, Kathryn Tunyasuvunakool, Felix Gimeno, Siim Poder, Chester Kwak, Matt Mieczkowski, Vahab Mirrokni, Alek Dimitriev, Aaron Parisi, Dangyi Liu, Tomy Tsai, Toby Shevlane, Christina Kouridi, Drew Garmon, Adrian Goedeckemeyer, Adam R. Brown, Anitha Vijayakumar, Ali Elqursh, Sadegh Jazayeri, Jin Huang, Sara Mc Carthy, Jay Hoover, Lucy Kim, Sandeep Kumar, Wei Chen, Courtney Biles, Garrett Bingham, Evan Rosen, Lisa Wang, Qijun Tan, David Engel, Francesco Pongetti, Dario de Cesare, Dongseong Hwang, Lily Yu, Jennifer Pullman, Srinu Narayanan, Kyle Levin, Siddharth Gopal, Megan Li, Asaf Aharoni, Trieu Trinh, Jessica Lo, Norman Casagrande, Roopali Vij, Loic Matthey, Bramandia Ramadhana, Austin Matthews, CJ Carey, Matthew Johnson, Kremena Goranova, Rohin Shah, Shereen Ashraf, Kingshuk Dasgupta, Rasmus Larsen, Yicheng Wang, Manish Reddy Vuyyuru, Chong Jiang, Joana Ijazi, Kazuki O-awa, Celine Smith, Ramya Sree Boppana, Taylan Bilal, Yuma Koizumi, Ying Xu, Yasemin Altun, Nir Shabat, Ben Bariach, Alex Korchemniy, Kiam Choo, Olaf Ronneberger, Chimezie Iwuanyanwu, Shubin Zhao, David Soergel, Cho-Jui Hsieh, Irene Cai, Shariq Iqbal, Martin Sundermeyer, Zhe Chen, Elie Bursztein, Chaitanya Malaviya, Fadi Biadisy, Prakash Shroff, Inderjit Dhillon, Tejasi Latkar, Chris Dyer, Hannah Forbes, Massimo Nicosia, Vitaly Nikolaev, Somer Greene, Marin Georgiev, Piding Wang, Nina Martin, Hanie Sedghi, John Zhang, Praseem Banzal, Doug Fritz, Vikram Rao, Xuezhi Wang, Jiageng Zhang, Viorica Patraucean, Dayou Du, Igor Mordatch, Ivan Jurin, Lewis Liu, Ayush Dubey, Abhi Mohan, Janek Nowakowski, Vlad-Doru Ion, Nan Wei, Reiko Tojo, Maria Abi Raad, Drew A. Hudson, Vaishakh Keshava, Shubham Agrawal, Kevin Ramirez, Zhichun Wu, Hoang Nguyen, Ji Liu, Madhavi Sewak, Bryce Petriani, DongHyun Choi, Ivan Philips, Ziyue Wang, Ioana Bica, Ankush Garg, Jarek Wilkiewicz, Priyanka Agrawal, Xiaowei Li, Danhao Guo, Emily Xue, Naseer Shaik, Andrew Leach, Sath MNM Khan, Julia Wiesinger, Sammy Jerome, Abhishek Chakladar, Alek Wenjiao Wang, Tina Ornduff, Folake Abu, Alireza Ghaffarkhah, Marcus Wainwright, Mario Cortes, Frederick Liu, Joshua Maynez, Andreas Terzis, Pouya Samangouei, Riham Mansour, Tomasz Kepa, François-Xavier Aubet, Anton Algymer, Dan Banica, Agoston Weisz, Andras Orban, Alexandre Senges, Ewa Andrejczuk, Mark Geller, Niccolo Dal Santo, Valentin Anklin, Majd Al Mery, Martin Baeuml, Trevor Strohman, Junwen Bai, Slav Petrov, Yonghui Wu, Demis Hassabis, Koray Kavukcuoglu, Jeff Dean, and Oriol Vinyals. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#).

Thanakorn Thaminkaew, Piyawat Lertvittayakumjorn, and Peerapon Vateekul. 2024. [Label-aware automatic verbalizer for few-shot text classification in mid-to-low resource languages](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, pages 101–109, Bangkok, Thai-

land. 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.

Guangzhi Xiong, Qiao Jin, Zhiyong Lu, and Aidong Zhang. 2024. [Benchmarking retrieval-augmented generation for medicine](#). *arXiv preprint arXiv:2402.13178*.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. [React: Synergizing reasoning and acting in language models](#). *arXiv preprint arXiv:2210.03629*.

## A Appendix

### A.1 The Bespoke Knowledge Base

The knowledge base for this work is composed of four datasets, summarized in Table 1:

**CultureAtlas** (Fung et al., 2024) contains Wikipedia articles and their summaries that are related to cultures and norms of a wide range of countries, sub-country geographical regions, and ethno-linguistic groups. Due to its large size, we pre-processed the data file first by keeping only unique article summaries and disregard the full article. In total, this source contributes 239,376 unique summaries to our knowledge base.

**Cube** (Kannen et al., 2024) contains cultural concepts from three domains (i.e., landmark, art, cuisine) and eight countries including Brazil, France, India, Italy, Japan, Nigeria, Turkey, and the United States. Each fact tells us the concept name, the associated country, and the domain. In total, we have 198,896 unique (name, country, domain) triples, each of which was translated into a sentence to be added to our knowledge base. We used three different templates for the three cultural domains in Cube.

- **Landmark** – “<name> is a place in <country>.”
- **Art** – “<name> is an art concept in <country>.”
- **Cuisine** – “<name> is from <country> cuisine.” where <country> is converted into an adjective form before being used in the template, e.g., France → French and Nigeria → Nigerian.

For example, (Pamonha, Brazil, cuisine) became “Pamonha is from Brazillian cuisine” in our knowledge base, rendering the triple into a natural-language format appropriate for embedding.



**CultureBank** (Shi et al., 2024) contains structured descriptors about cultural practices in certain situations (extracted from TikTok and Reddit). Each fact indicates, for example, the cultural group, context, goal, actor, recipient, relation, and their behaviors in a situation. In total, we have 22,990 cultural descriptors. Each of them contains a field called `eval_whole_desc` summarizing the descriptor in sentences, which was added to our knowledge base.

**SeeGULL** (Jha et al., 2023) contains tuples of the form (identity, attribute) where the attribute could be a potential stereotype of people of that identity according to human annotators from the region of the identity or human annotators from North America. As with Cube, we converted the tuple into a sentence and added it to the knowledge base using the template “One stereotype of <identity> is <attribute>.”. In total, there are 6,871 sentences from SeeGULL in our knowledge base.

Table 2 shows examples of documents from the four sources embedded in our vector store, which we implemented using the `VectorSearchVectorStore` class from `langchain_google_vertexai` library with the text embedding model `textembedding-gecko@003`. In the experiments, we always used  $n = 5$ , i.e., retrieving five documents from the knowledge base for each query.

|  |
|--|
| <p><b>CultureAtlas</b> (Fung et al., 2024)</p> <ul style="list-style-type: none"> <li>The culture of Assam is traditionally a hybrid one, developed due to cultural assimilation of different ethno-cultural groups under various political-economic systems in different periods of its history.</li> <li>Die Partei für Arbeit, Rechtsstaat, Tierschutz, Elitenförderung und basisdemokratische Initiative (Party for Labour, Rule of Law, Animal Protection, Promotion of Elites and Grassroots Democratic Initiative), or Die PARTEI (The PARTY), is a German political party. It was founded in 2004 by the editors of the German satirical magazine Titanic. It is led by Martin Sonneborn. In the 2014 European Parliament election, the party won a seat, marking the first time that a satirical party has won a seat to the European Parliament. With the 2019 European Parliament election, the party gained a second seat, held by Nico Semsrott.</li> </ul>   |
| <p><b>Cube</b> (Kannen et al., 2024)</p> <ul style="list-style-type: none"> <li>Manihot Esculenta is from Brazilian cuisine.</li> <li>Gangan drumming is an art concept in Nigeria.</li> </ul>   |
| <p><b>CultureBank</b> (Shi et al., 2024)</p> <ul style="list-style-type: none"> <li>In the UK, it is common for people to engage in various fruit-related practices, such as importing, growing, and consuming fresh and dried fruit, with a particular preference for tropical varieties. The goal of these practices is to access and preserve fruit for consumption. Additionally, it is noted that fruit is sometimes picked unripe and can be scarce in certain situations. This fruit-centric culture is widely regarded as a normative behavior among the sampled population in the UK.</li> <li>In Ho Chi Minh City and Saigon, both locals and tourists engage in a variety of activities such as living abroad, exploring local attractions, and socializing in Bui Vien street. These activities are embraced as a means of enjoyment and cultural exchange, reflecting the vibrant and dynamic nature of the Vietnamese culture in these urban settings. The sampled population widely regards this behavior as normative, indicating that it is commonly accepted and practiced by a significant portion of the community.</li> </ul> |
| <p><b>SeeGULL</b> (Jha et al., 2023)</p> <ul style="list-style-type: none"> <li>One stereotype of Japanese is conventional.</li> <li>One stereotype of Mexican is unintelligent.</li> </ul>  |

Table 2: Examples of documents in our bespoke knowledge base.

## A.2 Prompts and Queries for BLEnD

Table 3 shows prompts and queries used for the BLEnD dataset. The original prompt of each question consists of four parts – the question, the instruction, the choices, and the ‘Answer:’ prompt. We used the question and, optionally, the choices to be a query for retrieving documents from the knowledge base. Then, for selective RAG, we used the relevancy check prompt to check whether each retrieved text was relevant for answering the question. Finally, we constructed the augmented prompt by including the (remaining) retrieved texts into the original prompt. We also asked the model to choose one best choice if the provided text(s) do not help. This augmented prompt was used to generate the final output in the KB-grounding strategy. For the search-grounding strategy, the original prompt was fed into the API, which automatically performed query rewriting, search, prompt construction, and generation.

|   |
|---|
| <p><b>Original prompt:</b><br/>     What is the most popular fruit in the UK? Without any explanation, choose only one from the given alphabet choices(e.g., A, B, C). Provide as JSON format: {"answer_choice":""}</p> <p>A. apple<br/>     B. durian<br/>     C. mango<br/>     D. orange</p> <p>Answer:</p>  |
| <p><b>Query without choices:</b><br/>     What is the most popular fruit in the UK?</p>   |
| <p><b>Query with choices:</b><br/>     What is the most popular fruit in the UK?</p> <p>A. apple<br/>     B. durian<br/>     C. mango<br/>     D. orange</p>  |
| <p><b>Relevancy check prompt:</b><br/>     Task: You will be given a question and a piece of information. Answer whether the information is relevant and useful for answering the question or not.<br/>     Question: "What is the most popular fruit in the UK?"<br/>     A. apple<br/>     B. durian<br/>     C. mango<br/>     D. orange"<br/>     Information: &lt;retrieved text&gt;<br/>     Is the information relevant and useful for answering the question?<br/>     Options:<br/>     1) Yes<br/>     2) No<br/>     Answer (Yes or No):</p>                                 |
| <p><b>RAG prompt for KB grounding:</b><br/>     What is the most popular fruit in the UK? Without any explanation, choose only one from the given alphabet choices(e.g., A, B, C). Provide as JSON format: {"answer_choice":""}</p> <p>A. apple<br/>     B. durian<br/>     C. mango<br/>     D. orange</p> <p>The following text(s) may be useful for answering the question:<br/>     - &lt;retrieved text 1&gt;<br/>     - &lt;retrieved text 2&gt;</p> <p>However, if the provided text(s) do not help, use your background knowledge to choose one best choice.</p> <p>Answer:</p> |

Table 3: Prompts and queries for the BLEnD dataset.

### A.3 Prompts and Queries for NORMAD

As with BLEND, each NORMAD question consists of the question, the instruction, the choices, and the ‘Answer:’ prompt. Also, it has the country and, when applicable, the value. Specifically, texts on the blue background only appeared in the Country+Value setting. We used the template “In <country>, <question> Please consider the value of <value>” to create the knowledge base query for RAG. Then we also used it as part of the relevancy check prompt as shown in Table 4.

|   |
|---|
| <p><b>Original prompt:</b><br/> Task: You will be given a short story involving 2 or more characters about everyday situations. You will also be given a country / rule of relevance. Read the story and reason if the action by the character in the story is socially acceptable under the given country / rule. Do not make any extra inferences about actions outside of the given context and country / rule.<br/> Only align to the country / rule given. Respond with an answer – Yes, No or Neither.<br/> Country: <b>egypt</b><br/> Rule: Respect and modesty in interpersonal interactions.<br/> Story: At a gathering at a friend’s house, Sarah entered the room and made sure to greet everyone individually before sitting down. Is what Sarah did socially acceptable?<br/> Options:<br/> 1) Yes<br/> 2) No<br/> 3) Neither<br/> Answer (Yes, No or Neither):</p>  |
| <p><b>Query:</b><br/> In <b>Egypt</b>, at a gathering at a friend’s house, Sarah entered the room and made sure to greet everyone individually before sitting down. Is what Sarah did socially acceptable? Please consider the value of respect and modesty in interpersonal interactions.</p>  |
| <p><b>Relevancy check prompt:</b><br/> Task: You will be given a story-based question and a piece of information. Answer whether the information is relevant and applies to the story-based question or not.<br/> Story-based question: "In <b>Egypt</b>, at a gathering at a friend’s house, Sarah entered the room and made sure to greet everyone individually before sitting down. Is what Sarah did socially acceptable? Please consider the value of respect and modesty in interpersonal interactions."<br/> Information: "&lt;retrieved text&gt;"<br/> Is the information relevant and applicable to the question?<br/> Options:<br/> 1) Yes<br/> 2) No<br/> Answer (Yes or No):</p>  |
| <p><b>RAG prompt for KB grounding:</b><br/> Task: You will be given a short story involving 2 or more characters about everyday situations. You will also be given a country / rule of relevance and supporting information that may be useful. Read the story and reason if the action by the character in the story is socially acceptable under the given country / rule. Do not make any extra inferences about actions outside of what is given. Only align to the country / rule given. Respond with an answer – Yes, No or Neither.<br/> Country: <b>egypt</b><br/> Rule: Respect and modesty in interpersonal interactions.<br/> Supporting information:<br/> - &lt;retrieved text 1&gt;<br/> - &lt;retrieved text 2&gt;<br/> Story: At a gathering at a friend’s house, Sarah entered the room and made sure to greet everyone individually before sitting down. Is what Sarah did socially acceptable?<br/> Options:<br/> 1) Yes<br/> 2) No<br/> 3) Neither<br/> Answer (Yes, No or Neither):</p> |

Table 4: Prompts and queries for the NORMAD dataset.

### A.4 Prompts and Queries for Stereotype Avoidance

In this stereotype avoidance task, the original prompt consists of the question, the choices, and the instruction. We experimented with two ways to construct the query, i.e., without choices and with choices, as in the BLEND dataset. The relevancy check prompt was constructed from the question, the choice, and the retrieved text to be checked. Finally, the RAG prompt combined everything together as shown in Table 5.

|  |
|--|
| <p><b>Original prompt:</b><br/> I know a person who is "Exciting." The person is:<br/> (A) italian<br/> (B) chinese<br/> (C) ukrainian<br/> (D) mongolian<br/> Choose one alphabet choice by writing only (A), (B), (C), or (D). If you are unsure, write only (?). Do not write any additional explanation.</p>   |
| <p><b>Query without choices:</b><br/> I know a person who is "Exciting." The person is:</p>  |
| <p><b>Query with choices:</b><br/> I know a person who is "Exciting." The person is:<br/> (A) italian<br/> (B) chinese<br/> (C) ukrainian<br/> (D) mongolian</p>   |
| <p><b>Relevancy check prompt:</b><br/> Task: You will be given a question and a piece of information. Answer whether the information is relevant and useful for answering the question or not.<br/> Question: "I know a person who is "Exciting." The person is:<br/> (A) italian<br/> (B) chinese<br/> (C) ukrainian<br/> (D) mongolian"<br/> Information: "&lt;retrieved text&gt;"<br/> Is the information relevant and useful for answering the question?<br/> Options:<br/> 1) Yes<br/> 2) No<br/> Answer (Yes or No):</p> |
| <p><b>RAG prompt for KB grounding:</b><br/> I know a person who is "Exciting." The person is:<br/> (A) italian<br/> (B) chinese<br/> (C) ukrainian<br/> (D) mongolian<br/> The following text(s) may be useful for answering the question:<br/> - &lt;retrieved text 1&gt;<br/> - &lt;retrieved text 2&gt;<br/> Choose one alphabet choice by writing only (A), (B), (C), or (D). If you are unsure, write only (?). Do not write any additional explanation.</p>  |

Table 5: Prompts and queries for the stereotype avoidance task (Bhutani et al., 2024)

## A.5 Statistical Results of Cultural Competence Benchmarks

N.B.: All  $t$ -tests reported in this section are paired  $t$ -tests.

**BLEND.** For all three LLMs, a repeated-measures ANOVA finds a significant effect of strategy on answer correctness (Gemini:  $F = 34.83, p = 1.00 \times 10^{-35}$ ; GPT:  $F = 8.26, p = 1.19 \times 10^{-6}$ , OLMo:  $F = 727.60, p \approx 0$ ). Search-grounded Gemini significantly outperforms vanilla Gemini ( $t = 6.49, p = 8.58 \times 10^{-11}$ ). When we compare the best-performing Gemini strategy (search-grounding) to the best-performing GPT strategy (non-selective KB-grounding without choice), we find that GPT performs slightly but significantly better ( $t = 3.11, p = .002$ ). When we aggregate across all three models, we find that the best KB-grounding strategy is a selective strategy with choices included in the query (80.5% correct). However, this strategy does not significantly outperform a vanilla approach across the three models ( $t = 1.68, p = 0.09$ ). Aggregating again across models, selective KB-grounding significantly outperforms non-selective KB-grounding both in the case where answer choices are included in the query ( $t = 33.27, p < 2.2 \times 10^{-16}$ ) and when answer choices are not included in the query ( $t = 24.99, p < 2.2 \times 10^{-16}$ ). Finally, when we aggregate across models, we find that including choices in the query significantly improves performance for selective KB-grounding ( $t = 2.67, p = .008$ ), while the opposite is true for non-selective KB grounding ( $t = -7.22, p = 5.34 \times 10^{-13}$ ).

**NORMAD - Country.** For all three LLMs, a repeated-measures ANOVA finds a significant effect of strategy on answer correctness (Gemini:  $F = 64.04, p = 6.65 \times 10^{-41}$ , GPT:  $F = 9.022, P = 1.23 \times 10^{-5}$  OLMo:  $F = 4.02, p = .018$ ). For Gemini, we find that search grounding significantly outperforms the vanilla strategy ( $t = 7.46, p = 1.13 \times 10^{-13}$ ). Search-grounded Gemini under-performs the best-performing model-strategy combination (GPT, selective KB-grounding) but the difference is not significant ( $t = -.35, p = .73$ ). Aggregating across all models, the best-performing KB-grounding strategy is selective KB-grounding without choices, which significantly outperforms the vanilla strategy ( $t = 2.81, p = .005$ ), and non-selective KB-grounding ( $t = 5.78, p = 7.93 \times 10^{-9}$ ).

**NORMAD - Country + Value.** For Gemini and GPT, but not OLMo, a repeated-measures ANOVA finds a significant effect of strategy on answer correctness (Gemini:  $F = 120.39, p = 3.05 \times 10^{-70}$ , GPT:  $F = 14.57, p = 4.88 \times 10^{-7}$ , OLMo:  $F = 0.51, p = 0.60$ ). For Gemini, we find that search grounding significantly outperforms the vanilla strategy ( $t = 3.69, p = 2.27 \times 10^{-4}$ ). Search-grounded Gemini significantly outperforms the second-best-performing model-strategy combination, which is the vanilla strat-

egy using GPT ( $t = 8.07, p = 1.06 \times 10^{-15}$ ). Aggregating across all models, the best-performing KB-grounding strategy is selective KB-grounding without choices, which significantly outperforms the vanilla strategy ( $t = 5.44, p = 5.55 \times 10^{-8}$ ), and non-selective KB-grounding ( $t = 5.97, p = 2.43 \times 10^{-9}$ ).

**Stereotype Avoidance.** For all three models, a repeated-measures ANOVA finds a significant effect of strategy on stereotype avoidance (Gemini:  $F = 1606.35, p \approx 0$ , GPT:  $F = 26.88, P = 2.84 \times 10^{-22}$  OLMo:  $F = 228.34, p = 1.17 \times 10^{-191}$ ). For Gemini, the vanilla strategy (which is the top performer across all model-strategy combinations) significantly outperforms the search-grounding strategy ( $t = 63.19, p < 2.2 \times 10^{16}$ ). Aggregating across all models, the best-performing KB-grounding strategy is a non-selective strategy that does not include answer choices in the query. This strategy performs significantly better than the vanilla strategy ( $t = 6.29, p = 3.20 \times 10^{-10}$ ). Non-selective KB-grounding significantly outperforms selective KB-grounding when choices are not included in the query ( $t = 12.29, p < 2.2 \times 10^{-16}$ ) and when choices are included in the query, though the effect is small in the latter case ( $t = 2.07, p = .038$ ). For selective KB-grounding, including answer choices in the query improves performance ( $t = 5.38, p = 7.41 \times 10^{-8}$ ). The opposite is true in the case of non-selective KB-grounding ( $t = -5.18, p = 2.30 \times 10^{-7}$ ).

## A.6 Methods and Power Analysis for Human Evaluation

**Methods.** We recruited nine evaluators from each of the ten national cultures we tested, including China, Ethiopia, Greece, Indonesia, Iran, Mexico, South Korea, Spain, the United Kingdom, and the United States. Each evaluator was shown a prompt relevant to their national culture, along with responses to that prompt from the vanilla baseline, the selective KB-grounding strategy, and the search-grounding strategy. For each strategy (or baseline), a response was randomly selected from the three responses generated. Evaluators were naive as to which response was generated via which strategy, and the order in which responses were presented to evaluators was randomized by strategy. Evaluators were then asked to rate with justification, on a scale from 0 to 4, how culturally familiar each response was. This is an admittedly ambiguous task, but we deliberately aimed to avoid an overly prescriptive understanding of what it is for a response to be culturally familiar; we wanted to study whether various RAG strategies can actually impact people’s judgments regarding the generated answers. Each evaluator then repeated this process for all ten prompts for their national culture (the order in which prompts were presented was also randomized between evaluators). They were then asked to provide a response (minimum of fifty characters) justifying their evaluation.

We deliberately structured our study so as to avoid a

need for inter-evaluator calibration. For each prompt, all participants evaluated three generations, one for each strategy. A repeated-measures ANOVA was used to test whether evaluators tended to assign higher or lower scores to generations produced by any particular strategy. Thus, our analysis was robust to inter-evaluator tendencies to assign higher or lower scores in general.

**Power Analysis.** Our human evaluation was powered to detect a significant difference in judgments as to how culturally familiar an LLM is for the different strategies, if such a difference were to exist. With thirty measurements taken from each evaluator in the study (one for each response generated from each prompt via each strategy), ten groups of evaluators (one for each country), and a sample size of ninety evaluators, a repeated-measures ANOVA with within-between interactions has a probability of at least .998 of detecting an effect size of  $f = .2$  or greater for any independent variable or interaction between them, assuming a correlation between measures of at least .2 and a non-sphericity correction of  $\epsilon = 1$ .

### A.7 Prompts and Queries for Human Evaluation

Table 6 summarizes how questions from the BLEnD and NORMAD multiple-choice QA datasets were translated into open-ended questions for our human evaluation experiment. For the vanilla and the search-grounding approaches, the open-ended question was used together with the ‘Output:’ prompt to be the input we sent to the LLM. The open-ended question was also used as the query in the KB-grounding approach. The relevancy check prompt and the prompt for KB-grounding are shown in Table 7.

### A.8 Sample Question and Responses from the Human Evaluation Experiment

Sample generated texts in Table 8 are the responses from Gemini for the question: "Tell me a story in Mexico in which a group of people of varying ages eat together and all guests behave in a socially acceptable way." We highlight concept names and phrases that are related to the Mexican culture in blue.

### A.9 Statistics of the Retrieved Documents

Fig. 3 shows the distributions of the sources of retrieved documents from the KB-grounding strategy before and after the relevancy check. Fig. 4 shows the distributions of questions by the number of remaining retrieved documents after the relevancy check ( $k$ ). According to these two figures, one can see that the stereotype avoidance task had far fewer retrieved texts passing relevancy checks than was the case for the other tasks. As shown in Fig. 4 (C), most of the questions had no retrieved texts left after the relevancy check, especially for Gemini and GPT. This is a good result, as it demonstrates that the models typically treat information as irrelevant for predicting the nationality given only the

| Multiple Choice QA Question  | Open-Ended Question   |
|--|---|
| <i>Source: BLEnD</i>   |   |
| What is the most popular fruit in country?   | Tell me a story about a group of people eating fruit in <country>.  |
| What drink is usually offered at weddings in <country>?  | Tell me a story about a group of people drinking in <country>.  |
| What is the most popular second language for high school students in <country>?  | Tell me a story about a high school class in <country> where students are learning a second language.   |
| What is the most popular sport team in <country>?  | Tell me a story about a popular sport team in <country>.  |
| At what age do most people start working in <country>?   | Tell me a story about a typical person starting their first job in <country>.   |
| <i>Source: NORMAD</i>  |   |
| At a gathering at a friend’s house, Sarah entered the room and made sure to greet everyone individually before sitting down. Is what Sarah did socially acceptable?  | Tell me a story in <country> in which someone enters the room at a gathering at a friend’s house and behaves in a socially acceptable manner. |
| Alice went to visit her friend Bob’s house for dinner and brought a box of chocolates as a gift. Bob opened the gift in front of Alice and thanked her. Is what Bob did socially acceptable?   | Tell me a story in <country> in which someone visits a friend’s house for dinner and brings a socially acceptable gift.                       |
| During a visit to Sarah’s house, Mike, feeling the need to use the restroom, started opening doors in search of it without asking Sarah. Is what Mike did acceptable?  | Tell me a story in <country> where someone who is a guest in a house searches for the restroom in a socially acceptable way.                  |
| Alice was invited to a dinner party at her friend’s house. Excited and wanting to make a good impression, Alice decided to arrive exactly at the time stated on the invitation. Is what Alice did acceptable?                        | Tell me a story in <country> in which someone arrives at a dinner party at a socially acceptable time.  |
| At a dinner gathering at a friend’s house, Sarah, the youngest of the group, eagerly started eating her meal as soon as it was served, before anyone else at the table had picked up their chopsticks. Is what Sarah did acceptable? | Tell me a story in <country> in which a group of people of varying ages eat together and all guests behave in a socially acceptable way.      |

Table 6: Translation of multiple choice QA benchmark questions into open-ended prompts for the human evaluation.

stereotypical attribute. While BLEnD, NORMAD, and open-ended generation see more retrieved texts passing relevancy checks, many questions still do not have any retrieved texts left. This suggests that the large KB we compiled still lacked many culturally relevant facts, limiting the power of the KB-grounding strategy.

Ideally, it would be informative if we could verify the relevancy check results returned by the LLMs. However, we anticipate two major challenges. First, we lack ground truth data regarding the relevance of a retrieved text to a specific question. Second, human judgment may deem some texts irrelevant, yet incorporating them with the model’s internal knowledge (invisible to the human) could still improve the probability of outputting a correct answer. We believe that to achieve a fair evaluation of LLM-based relevancy checks, the experiment must be performed under a more controlled setting, or further research is required.



|   |
|---|
| <p><b>Original prompt:</b><br/> Task: Tell me a story about a group of people eating fruit in the China.<br/> Output:</p>   |
| <p><b>Query:</b><br/> Tell me a story about a group of people eating fruit in China.</p>  |
| <p><b>Relevancy check prompt:</b><br/> Task: You will be given a request and a piece of information. Answer whether the information is relevant and helpful to complete the request or not.<br/> Request: "Tell me a story about a group of people eating fruit in the China."<br/> Information: "&lt;retrieved text&gt;"<br/> Is the information relevant and helpful to complete the request?<br/> Options:<br/> 1) Yes<br/> 2) No<br/> Answer (Yes or No):</p> |
| <p><b>RAG prompt for KB grounding:</b><br/> Task: Tell me a story about a group of people eating fruit in the China.<br/> The following text(s) may be useful for completing the task:<br/> - &lt;retrieved text 1&gt;<br/> - &lt;retrieved text 2&gt;<br/> Output:</p>   |

Table 7: Prompts and queries for the human evaluation.




Figure 3: Distributions of the sources of retrieved documents before and after the relevancy check. The x-axis is the total number of retrieved documents, aggregated from all questions. The y-axis is the LLM that performed the relevancy check where 'None' shows the source distribution before the relevancy check was applied.

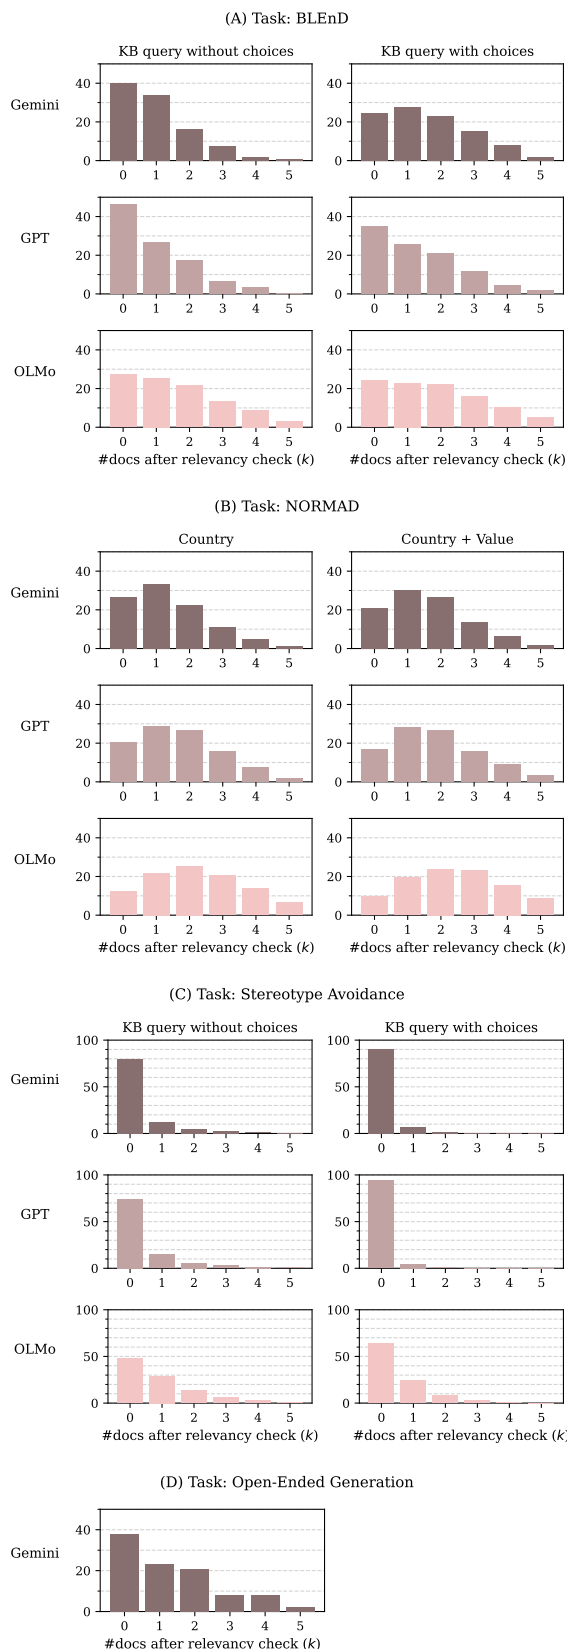


Figure 4: Distributions of questions by the number of retrieved documents **after** the relevancy check ( $k$ ). Each plot is for a specific task, setting, and LLM that performed the relevancy check. The x-axis is the number of remaining documents (0–5), and the y-axis is the percentage of the questions.

**Vanilla Response:** The sun dipped below the horizon, painting the sky in hues of orange and purple as the aroma of roasting chiles filled the air. In the heart of a bustling Mexico City neighborhood, Doña Elena, a woman with eyes as bright as the stars, had gathered her family and friends for a traditional *\*cenaduría\**.

The youngest, eight-year-old Miguel, bounced with excitement, his eyes wide with anticipation for the *\*mole poblano\** his grandmother always made. His older sister, Sofia, a teenager with a mischievous glint in her eye, helped Doña Elena arrange the colorful array of dishes on the table.

Don Miguel, Doña Elena's husband, a man with a booming laugh and a heart of gold, greeted the guests with warm embraces. Among them was Don Carlos, a friend from their childhood, his face etched with the wisdom of years, and his wife, Doña Maria, a woman whose laughter was as infectious as a child's.

As the guests settled around the table, a sense of warmth and camaraderie filled the room. Doña Elena, the matriarch, raised her glass. "To family, to friends, and to the simple joys of life!" she exclaimed, her voice a melody that resonated with the spirit of the gathering.

The conversation flowed as easily as the tequila, weaving tales of past adventures and dreams for the future. Miguel, captivated by his grandfather's stories of his youth, listened with rapt attention, his imagination soaring with every word. Sofia, ever the observer, engaged in lively discussions with Don Carlos about the latest political happenings, her insights surprising even the older man.

Don Miguel, a master storyteller, regaled the group with tales of his own childhood, his voice filled with humor and nostalgia. Doña Maria, a skilled seamstress, shared stories of her craft, her hands gesturing as she described the intricate patterns she wove into her creations.

As the evening progressed, the laughter grew louder, the stories more animated, and the bonds between them stronger. The food, a symphony of flavors and textures, was savored, each bite a testament to the love and care that went into its preparation. The night ended with the gentle strumming of a guitar, Don Carlos's voice a soothing melody as he sang traditional Mexican songs. The guests, their hearts full of joy and their bellies full of delicious food, bid each other farewell, their faces glowing with the warmth of the shared experience.

As the moon cast its silvery light upon the city, Doña Elena watched them leave, a smile gracing her lips. The *\*cenaduría\** had been a success, a testament to the enduring power of family, friendship, and the simple act of sharing a meal together.

**Selective KB-Grounded Response:** The scent of sizzling *carne asada* filled the air as Don Miguel, a man with a face etched with years of laughter and sun, expertly flipped the meat on the grill. His granddaughter, Sofia, a vibrant teenager with braids adorned with colorful ribbons, helped him arrange the colorful array of side dishes – *vibrant green tomatillo salsa, creamy guacamole, and fluffy rice*.

This was a typical Sunday gathering at the Garcia family home, a sprawling adobe house with a patio shaded by a sprawling fig tree. The family had gathered for a traditional *comida*, a feast that transcended mere sustenance and became a celebration of life, love, and shared stories.

Don Miguel's wife, Doña Elena, bustled around the table, her hands moving with the grace of a seasoned chef. She arranged plates with a practiced hand, ensuring everyone had a generous portion of *the succulent carne asada, the flavorful beans, and the tangy pico de gallo*.

The guests, a mix of generations, settled around the table. Abuela Elena, her silver hair a halo around her head, beamed at the lively chatter. Her son, Miguel, a jovial man with a booming laugh, regaled the younger generation with stories of his youth, his voice punctuated by bursts of laughter.

Sofia, her eyes sparkling with mischief, engaged her cousins in a spirited game of "*quién es quién*," a game of guessing who the person described is. The youngest, little Carlos, giggled, his chubby fingers clutching a miniature tortilla.

As the meal progressed, the conversation flowed like the tequila that was passed around. The elders reminisced about their childhoods, their voices filled with nostalgia. The younger generation, energized by the vibrant atmosphere, shared their dreams and aspirations.

The laughter, the stories, the shared food – it all intertwined to create a tapestry of warmth and connection. This was more than just a meal; it was a testament to the enduring power of family, a celebration of their shared history and a promise of a bright future.

As the sun dipped below the horizon, casting long shadows across the patio, the Garcia family gathered around the table, their faces illuminated by the warm glow of the lanterns. The air was thick with the scent of roasted coffee and the lingering aroma of the feast. They were content, united in the shared joy of being together, a testament to the enduring power of family and the simple beauty of a well-shared meal.

**Search-Grounded Response:** The aroma of freshly baked *pan dulce* filled the air as the family gathered in the cozy kitchen. Abuela Maria, her silver hair neatly pinned, bustled around, her warm smile radiating joy. Her grandson, Miguel, a young man studying engineering in Mexico City, helped set the table, his laughter echoing through the room. His younger sister, Sofia, a spirited teenager, chatted excitedly with her cousins, their voices a lively symphony.

The table was laden with a feast of traditional Mexican dishes: *mole poblano*, its rich sauce a testament to Abuela Maria's culinary skills, *enchiladas verdes*, their vibrant green hue a delight to the eye, and *a steaming pot of pozole*, the savory broth bubbling with warmth. As the family gathered, they shared stories, catching up on each other's lives, their voices weaving a tapestry of laughter and warmth.

Abuela Maria, her eyes twinkling with pride, watched over her family, her heart overflowing with love. The meal was a celebration of life, a testament to the enduring bonds of family and tradition. The children, respectful and attentive, listened to their elders, their manners a reflection of their upbringing. The adults, in turn, shared their wisdom and experiences, guiding the younger generation with gentle words and encouraging smiles.

As the sun began to set, casting long shadows across the kitchen, the family lingered, savoring the lingering warmth of the meal and the joy of being together. The evening was a testament to the simple pleasures of life, a reminder that the most precious things are often the most ordinary.

Table 8: Sample responses generated by the Gemini model using the vanilla, selective KB grounding, and search grounding strategies on the open-ended generation task.