

Information Asymmetry Across Language Varieties: A Case Study on Cantonese-Mandarin and Bavarian-German QA

Renhao Pei^{*✉}, Siyao Peng^{*▲}, Verena Blaschke[▲], Robert Litschko[▲], Barbara Plank[▲]

▲MaiNLP, Center for Information and Language Processing, LMU Munich, Germany

✉Munich Center for Machine Learning (MCML), Germany

✉ELLIS Institute Finland ✉University of Turku, Finland

renhaopei@gmail.com loganpeng1992@gmail.com bplank@cis.lmu.de

Abstract

Large Language Models (LLMs) are becoming a common way for humans to seek knowledge, yet their coverage and reliability vary widely. Especially for *local language varieties*, there are large asymmetries, e.g., information in local Wikipedia that is absent from the standard variant. However, little is known about how well LLMs perform under such *information asymmetry*, especially on closely related languages. We manually construct a novel challenge question-answering (QA) dataset that captures knowledge conveyed on a local Wikipedia page, which is absent from their higher-resource counterparts—covering Mandarin Chinese vs. Cantonese and German vs. Bavarian. Our experiments show that LLMs fail to answer questions about information only in local editions of Wikipedia. Providing context from lead sections substantially improves performance, with further gains possible via translation. Our topical, geographic annotations, and stratified evaluations reveal the usefulness of local Wikipedia editions as sources of both regional and global information. These findings raise critical questions about inclusivity and cultural coverage of LLMs.

Keywords: question answering, language varieties, information asymmetry, large language models

1. Introduction

Large Language Models (LLMs) are becoming a commonly used channel through which humans find and learn new information (Brachman et al., 2025; Elsey et al., 2025; Yang et al., 2025b). Another information source is Wikipedia, which includes editions in over 300 languages¹ and any article can be created or edited independently (Lewoniewski et al., 2017), so that information about the same topic could vary largely across languages (Callahan and Herring, 2011; Bao et al., 2012; Wang et al., 2025). Till now, far less is understood about how LLMs handle facts relevant to specific cultural or linguistic communities that are less visible in global discourse (Myung et al., 2024).

Cross-lingual retrieval for information-seeking questions remains challenging for LLMs. Answer retrieval during post-training remains the main bottleneck even when provided with contextual information (Gao et al., 2025). Such challenges are particularly pronounced in open-book question answering (QA), where additional external information is provided, in contrast to closed-book QA, which relies on knowledge already in the model’s training (Yen et al., 2023). LLMs also struggle to share knowledge across languages, i.e., when the question is proposed in a language different from the context language (Goldman et al., 2025).

In this work, we address cross-lingual context retrieval on regional knowledge. Specifically, we focus on how well LLMs can answer questions drawn from knowledge in local-language Wikipedia pages but which are absent from standard-language pages. This is similar to the *Information Asymmetry* analyzed in Roy et al. (2020, 2022). However, rather than focusing on the number of articles covered and the length of articles in English vs. non-English Wikipedia versions, we frame it as a QA task, seeking local information that is absent from the standard Wikipedia editions. This aligns with Asai et al. (2021), who observe that questions about entities from other cultures may lack answers in the questioner’s language due to cultural bias, and with the notion of a “content gap” introduced by Duderstadt (2025).

To address this gap, we conduct a case study of two pairs of higher-resource (*standard*) versus lower-resource (*local*) language varieties spoken in overlapping regions: Mandarin Chinese versus Cantonese, and German versus Bavarian. These pairs allow us to probe whether LLMs capture regional facts beyond what is available in dominant, higher-resource variants. We experiment with closed-book QA, i.e., relying on LLMs’ internal parametric knowledge, and open-book QA, i.e., explicitly providing the extracted local information. We focus on Wikipedia lead sections, which provide concise yet systematically curated knowledge that are length-friendly to test LLMs’ capability in open-book QA.

We construct a new evaluation dataset, WiLOVA-

*Equal contribution.

¹https://meta.wikimedia.org/wiki/List_of_Wikipedias

QA, consisting of facts that appear in local lead sections but are absent in their standard counterparts. Our dataset includes 294 lead section QA pairs for Cantonese–Mandarin and 178 for Bavarian–German. Furthermore, we annotate a smaller set of document-level QA pairs for which the information is extracted from the full local Wikipedia article beyond the local lead section, yet absent from the aligned standard Wikipedia page. These document-level annotations introduce an additional evaluation scenario for seeking local information beyond the provided context.

Our experiments yield four key findings. First, LLMs consistently fail to answer questions derived from local Wikipedia lead sections in closed-book QA, suggesting that such knowledge is systematically underrepresented in models. Second, providing lead sections from local Wikipedia editions as context substantially improves performance, indicating that LLMs can reason over that knowledge when it is explicitly supplied. Translating these lead sections into the standard variant, i.e., from Cantonese to Mandarin, further boosts performance, highlighting the importance of context language for knowledge integration. Third, we evaluate on a cross-lingual dataset, ECLeKTic (Goldman et al., 2025), and demonstrate that QA performance on local language varieties still lags behind when compared to seeking information in widely spoken languages. Lastly, we manually categorize WiLoVA-QA articles to assess their topic and regional relevance and their effects on model performance. Better performance is achieved on articles discussing more regional topics than on others. Our results corroborate earlier findings (Roy et al., 2020, 2022; Wang et al., 2025) that local Wikipedia editions are not only useful as sources of regional knowledge but can also provide additional information on more region-independent topics. We release our dataset and code to facilitate future research on cultural representation and coverage in LLMs.²

2. Related Work

Earlier studies on BERT-based models have shown that these LMs encode surprising amounts of relational knowledge, facilitating open-domain QA even without fine-tuning (Petroni et al., 2019; Jiang et al., 2020; Roberts et al., 2020; Kassner and Schütze, 2020). With LLMs, information is typically accessed from model-internal parametric knowledge (Mallen et al., 2023) acquired during pre-training, or from externally provided knowledge in the case of retrieval-augmented LLMs. LLMs degrade markedly when not prompted in English (Zhang et al., 2023). Moreover, multilingual LLMs

²Available at <https://github.com/mainlp/WiLoVa-QA>

prefer to retrieve information in the same language as the query, and when no such information exists, they rely on high-resource languages (Myung et al., 2024; Sharma et al., 2025).

This is compounded by **information asymmetry**, i.e., not all content is equally present in volume and nature on the web (Roy et al., 2022). For example, prior work has shown that locally famous people may be described in a more favorable way in their own Wikipedia language (Kolbitsch and Maurer, 2006; Callahan and Herring, 2011), or that some aspects are not discussed altogether in some language editions (Samir et al., 2024; Wang et al., 2025). Collaborative behaviors can also vary across different language editions of Wikipedia (Bipat et al., 2018). As such, it can impact the values and knowledge that LLMs learn during pretraining.

Similarly, some **regional or local information** is often found **only** in local editions of Wikipedia. Roy et al. (2020, 2022) compare information coverage of English Wikipedia with eight other languages and show English articles often miss out on details in other Wikipedia editions, and that almost 50% of non-English articles lack English correspondence. Ashrafimoghari (2023) also observed such “information gap” between language editions. This raises the question of whether LLMs reflect the same information asymmetry as on the web.

Many multilingual QA datasets are created to address this language and cultural bias, including TyDi QA (Clark et al., 2020), XOR QA (Asai et al., 2021), INCLUDE (Romanou et al., 2025), BLEND (Myung et al., 2024), and CVQA (Romero et al., 2025). Faisal et al. (2021) introduce a spoken dialect QA benchmark and show large performance drops for certain dialects. Closest to our work, Goldman et al. (2025) create the ECLeKTic dataset by selecting Wikipedia articles that exist **only** in one language edition and prompting Gemini to generate QA pairs. Model performance on ECLeKTic degrades when the question is translated and proposed to LLMs in another language.

Our case study addresses slightly more complex situations in which Wikipedia pages exist in both local and standard language varieties, with more information on the local side. We investigate LLMs’ asymmetric parametric knowledge gaps and their interaction with externally provided context.

3. WiLoVA-QA Dataset

We create WiLoVA-QA (Wikipedia Local Variety Question Answering), a small dataset to assess LLMs’ capability in retrieving information that appears on a Wikipedia page in the **local language variety** but is absent from its **standard variety** counterpart. We discuss our target language varieties (§3.1), Wikipedia data preprocessing (§3.2),

<p>Lead section (Cantonese): 南戲係中國戲曲史上最成熟文藝戲劇，北宋末年至明朝初年（12世紀~14世紀）中國東南沿海流行，為區別同時代「北曲雜劇」，後人稱之為「南曲戲文」、「南戲」或「戲文」。南戲存在，令中國古代戲曲同古希臘戲劇同古印度戲劇並列為世界三大古代戲劇體系。 <i>Nanxi is the earliest mature literary and artistic drama in the history of Chinese drama. It was popular in the southeastern coast of China from the end of the Northern Song Dynasty to the early Ming Dynasty (12th century to the 14th century). The existence of Nanxi has placed ancient Chinese drama alongside ancient Greek drama and ancient Indian drama as the three largest ancient theater systems in the world.</i></p> <p>Lead section (Mandarin): 南戲是中國戲曲史上最成熟戲劇類型，元朝流行於中國南方，為區別同時代的「北曲雜劇」，後人稱之為南曲戲文、溫州南戲或戲文。 <i>Nanxi is the earliest mature type of drama in the history of Chinese drama. It was popular in southern China during the Yuan Dynasty. To distinguish it from the contemporary "Beiqu Zaju", later generations called it Nanqu Xiwen, Wenzhou Nanxi, or Xiwen.</i></p> <p>Question 1: 南戲北宋末年至明朝初年（12世紀~14世紀）在中國的什麼地方流行？ <i>Where in China was Nanxi popular from the end of the Northern Song Dynasty to the beginning of the Ming Dynasty (12th to 14th century)?</i> Answer 1: 東南沿海。 <i>Southeast coast.</i></p> <p>Question 2: 中國古代戲曲和什麼並列為世界三大古代戲劇體系？ <i>Alongside what is Chinese ancient drama regarded as one of the world's three major ancient dramas?</i> Answer 2: 古希臘戲劇和古印度戲劇。 <i>Ancient Greek drama and ancient Indian drama.</i></p>
<p>Lead section (Bavarian): Es Hiataamadl is a traditionella Voiksdanz, wo in Östareich und Bayern weit vabroadd is. En Danz gibts in vaschiedna Variantn. <i>Da Nama kimmt vom Danzliadl "Koa Hiataamadl mog i net". Da Danz hod de sejm Wuazin wia da Strohschneida, wo iwa ganz Eiropa vabroadd wor. The Hiataamadl is a traditional folk dance widespread throughout Austria and Bavaria. There are various variations of the dance. The name comes from the dance tune "Koa Hiataamadl mog i net (I don't like shepherd girls)." The dance has the same roots as the Strohschneider (straw cutter), which was widespread throughout Europe.</i></p> <p>Lead section (German): Hiataamadl ist der Name eines alpenländischen Volkstanzes, der sich in Österreich und Bayern großer Beliebtheit erfreut. <i>Hiataamadl is the name of an Alpine folk dance that is very popular in Austria and Bavaria.</i></p> <p>Question 1: Woher kommt der Name des Tanzes Hiataamadl? <i>Where does the name of the dance Hiataamadl come from?</i> Answer 1: Der Name des Tanzes kommt vom Tanzlied "Koa Hiataamadl mog i net". <i>The name of the dance comes from the dance tune "Koa Hiataamadl mog i net".</i></p> <p>Question 2: Aus welchem Tanz hat sich das Hiataamadl entwickelt? <i>From which dance did the Hiataamadl develop?</i></p> <p>Answer 2: Das Hiataamadl hat dieselben Wurzeln wie der Strohschneider, der in ganz Europa verbreitet war. <i>The Hiataamadl has the same roots as the Strohschneider, which was widespread throughout Europe.</i></p>

Table 1: Sample WiLoVA-QA annotations for `yue-cmn` page “南戏” and `bar-deu` page “Hiataamadl.”

and question-answer annotations (§3.3). We present detailed dataset statistics in Appendix A.

3.1. Target Language Varieties

We focus on two *local* versus *standard* language pairs:³ Cantonese (ISO 639-3: `yue`) vs. Mandarin Chinese (`cmn`) and Bavarian (`bar`) vs. German (`deu`). Both pairs are spoken in overlapping regions, where the standard language prevails in formal communication, while the local variety serves as a marker of regional and cultural identity.

Cantonese is a Sinitic language spoken primarily in the Guangdong province, Hong Kong, and Macau, as well as in diaspora communities outside of China. Bavarian is a group of non-standardized Upper German dialects spoken in southeastern Germany (Bavaria), most of Austria, and in northern Italy (South Tyrol). Though Cantonese and Bavarian have over 10 million speakers and are among the top 100 most spoken languages (Eberhard et al., 2025), their digital presence is relatively sparse. Joshi et al. (2020) labeled Cantonese and Bavarian as “1 - The Scraping-Bys” (i.e., with some amount of unlabeled data) in a 0-5 scale for language resource richness (0 the rarest and 5 the richest), whereas Mandarin and German are both “5 - The Winners” (i.e., quintessential resource-rich languages). For example, there are 10 times more pages on the Mandarin Chinese Wikipedia (1.5M) than on the Cantonese Wikipedia (146K), and 100 times more on the German Wikipedia (3.0M) than

on the Bavarian Wikipedia (27K). Though underrepresented, these local pages still provide information that does not appear in standard variety pages, as shown in Table 1. This asymmetry makes Cantonese and Bavarian instructive test cases for studying how knowledge in regional Wikipedia editions is represented or omitted in LLMs.

3.2. Preprocessing Wikipedia Pages

Aligning Wikipedia pages We follow Litschko et al. (2025b) to extract and align Wikipedia pages from both the local and standard varieties. We obtain 111.6K and 24.6K aligned Wikipedia pages for Cantonese-Mandarin (`yue-cmn`) and Bavarian-German (`bar-deu`) and covering the majority of the local pages, 76.2% (111.6K/146.3K) of `yue` and 90.4% (24.6K/27.2K) of `bar`.

Local-heavy filtering Since our study seeks to manually annotate a small portion of information uniquely present on local pages, we further filter to local-heavy pairs. Namely, we only keep Wikipedia pairs for which both the *lead section* (typically the introductory paragraphs)⁴ and the whole document in the local language edition are longer than those of the corresponding page in the standard language. This resembles the `diff+` setting in Roy et al. (2020), where non-English articles exceed their English counterparts in length. We used the Wikipedia API⁵ to extract these lead sections (`page.summary`) and texts (`page.text`). After filtering, only 5.38% (6.2K/111.6K) and 4.07%

³For the sake of simplicity, we use the terms ‘local’ vs. ‘standard’. However, this terminology is somewhat flawed, as Cantonese is arguably a standardized language, especially in Hong Kong. However, speaker attitudes towards the status of Cantonese as a language (relative to Mandarin) vary greatly (Groves, 2010).

⁴https://en.wikipedia.org/wiki/Wikipedia:Manual_of_Style/Lead_section

⁵<https://pypi.org/project/Wikipedia-API/>.

(1.0K/24.6K) of `yue-cmn` and `bar-deu` pairs remained.

This simple filtering has some limitations. Our preprocessing and annotation only ensure that the information on the local-language page is absent from the *paired* standard page. It remains unknown whether such information is attested on other Wikipedia pages of the standard language, in resources outside of Wikipedia, or even in other languages beyond our study. Although other external sources of information may exist, our experiments and results clearly show that such information from these local Wikipedia pages remains difficult for LLMs to recall.

Removing direct translations and structural mismatches

Before our annotations, we noticed that some page pairs are not ideal for our information asymmetry annotation, even after our local-heavy filtering. These include cases that are (i) *translated*, if the local page seems to be a direct translation of the standard page (i.e., containing almost identical information and having a similar structure), regardless of whether the translated local variety sounds natural or superficial, and (ii) *structural mismatch* if either the local page or the standard page is a disambiguation page or a page without a lead section, lacking sufficient content for our information-asymmetry annotation. We ask annotators to exclude these two cases from QA annotation. We found that more than half of inspected `bar-deu` pages, 52.5% (261/497), are nearly direct translations, whereas only 26.6% (61/229) for `yue-cmn` pages. The proportion of structural mismatch pages is also higher in `bar-deu` (28.6%) than `yue-cmn` (10.5%). As a result, only 18.9% (94/497) of `bar-deu` pairs contain distinctive content in the local page for QA annotation; comparatively, more than half of `yue-cmn` pairs, 59.8% (137/229), deliver such information.

3.3. Question-Answering Annotation

We frame our local-heavy information-asymmetry annotation as a question-answering (QA) task. For each filtered Wikipedia pair, we ask annotators to compare both editions and identify information present only on the local page but absent from the standard. 137 `yue-cmn` and 94 `bar-deu` article pairs were annotated.

We conduct QA annotations at two levels. The main contributions are the *lead section* annotations. Wikipedia lead sections typically range from a few dozen to a few hundred tokens (much shorter than full Wikipedia documents), making it feasible for annotators to compare pages from two language varieties and extract information present only in the local one. Because our study evaluates LLMs on

answering such locally specific information—and compares performance with and without reference contexts in the prompt—using concise lead sections also keeps the experiments tractable. We additionally annotate a smaller number of *document-level* QAs to assess LLMs’ capability to seek out-of-context information, which could be useful for shortcut learning.

Lead section QA annotations For each pair of lead sections, we ask annotators to identify 2-3 pieces of information that are present only on the local side and missing on the standard side. We follow [Ogundepo et al. \(2023\)](#) in encouraging annotators to pursue their natural curiosity in discovering information asymmetry between versions. Annotators are instructed to (i) formulate a *question* in the standard language that addresses the information, (ii) *answer* the question in both the standard language and the local variety, and (iii) *highlight* texts on the local section from which the answer was drawn. We ask annotators to create questions that (a) are as diverse as possible, (b) can be answered with a short phrase, and (c) make sense as stand-alone questions. For example, we ask annotators to avoid frequently composing questions that ask for the year/date of an event or require multi-hop reasoning. WiLOVA-QA includes 294 `yue-cmn` and 178 `bar-deu` lead section QA annotations.

Verifying information asymmetry in the full document

For each lead section QA annotation, we ask annotators to verify whether the answer can be found in the full document of the standard page via a quick string search. If an answer is found, the supporting sentence(s) are copied into our annotation. Among 294 `yue-cmn` and 178 `bar-deu` lead section QA annotations, the majority, 200 (68.0%) and 100 (56.2%), **cannot** be answered even using the full content of the standard page.

Document-level QA annotations We supplement the WiLOVA-QA dataset with 80 and 46 document-level QA annotations for `yue-cmn` and `bar-deu`. We ask annotators to skim through the whole Wikipedia page of the local and standard language varieties and find one piece of information from the local page that is (a) not on the standard page and (b) not already annotated at the lead section level. Document-level QA annotation is conducted only when some additional information from the local page can be observed within 10-15 minutes; thus, it is not annotated on every Wikipedia page pair.

Annotators We recruited two master’s students in computational linguistics, one fluent in Cantonese and Mandarin and one fluent in Bavarian

and German. Two annotators were hired and compensated at national salary rates for 50 hours of annotation per language. Both annotators report that annotating each lead-section QA pair takes 5-10 minutes and that annotating each document-level QA takes 10-15 minutes, mostly due to reading the lead section and skimming both pages of the document. Because a larger portion of *bar-deu* pages were translations or structurally mismatched, more documents had to be screened and discarded, resulting in fewer QA annotations than for *yue-cmn*.

Quality Control Since the annotators were asked to extract a few pieces of information they found prominent from paired Wikipedia pages, their QA annotations remain non-exhaustive and open-ended. This makes traditional overlap-based inter-annotator agreement (IAA) inappropriate to evaluate the quality of WiLOVA-QA.

We follow [Ogundepo et al. \(2023\)](#) and ask two project coordinators, one native Mandarin speaker and one native German, to review the elicited questions and answers that are written in Mandarin or in German. Our quality check includes three aspects: (a) *grammaticality*: if any of the *question* or *answer* is ungrammatical or non-fluent; (b) *stand-alone validity*: if a *question* cannot function as an independent, context-free information-seeking question; and (c) *answer appropriateness*: if an answer does not directly and fully address its paired question when considered in isolation. Only a few questions and answers were flagged and rewritten on the entire dataset, demonstrating the overall high quality of our original QA annotations.

4. Experiments

We evaluate state-of-the-art LLMs’ ability to answer information-seeking questions whose answers are absent from the Wikipedia page in the prompting languages, namely Mandarin Chinese and German.

Models We evaluate six LLMs from three families: Llama3.1-8B/70B ([Grattafiori et al., 2024](#)), Qwen2.5-7B/72B ([Yang et al., 2025a](#)), and gpt-oss-20B/120B ([Agarwal et al., 2025](#)). We use vLLM ([Kwon et al., 2023](#)) as the inference backend for efficient LLM decoding. We set a low temperature of 0.3 to encourage more deterministic and consistent responses for the QA task. Models automatically choose the compute precision based on their respective settings (`dtype='auto'`), with a maximum context length of 10K tokens and a maximum output length of 5K tokens.

Datasets We use our lead section and document QA annotations to evaluate LLMs’ ability to answer

information-seeking questions. We also evaluate on ECLeKTic ([Goldman et al., 2025](#)) to compare performance when retrieving information from another standard language (e.g., French) versus a local variety (i.e., Cantonese or Bavarian).

Unlike our manual QA annotations, the ECLeKTic dataset generated knowledge-seeking questions in 12 source languages using Gemini ([Anil et al., 2025](#)), for which the answer appears in a relevant Wikipedia article in the source language, but the corresponding Wikipedia pages in the other 11 languages are non-existent. Questions and answers are translated into the other 11 languages to evaluate LLMs in cross-lingual knowledge transfer. To compare with our WiLOVA-QA results, we assess our LLM selection on a subset of 333 items sourced from the 10 languages, excluding Mandarin Chinese (35 items) and German (16 items). It is also worth noting that the extracted context in ECLeKTic includes the first 10 sentences of the Wikipedia pages, which differ slightly from the lead sections in WiLOVA-QA. Additionally, the answers in ECLeKTic are much shorter (on average 8.3 characters for *cmn* and 21.0 for *deu*) than WiLOVA-QA (21.7 for *cmn* and 83.9 for *deu*).

Context scenarios We prompt LLMs to answer information-seeking questions in WiLOVA-QA when given different context scenarios: (1a) *question-only*: asking the question without any additional context information; (1b) *+standard*: providing lead section text from the Wikipedia article in the standard language as context; (1c) *+local*: providing lead section text from the Wikipedia article in the local language; (1d) *+standard+local*: providing lead section texts from both the standard and the local from Wikipedia articles. Since the QA annotations are sourced from the Wikipedia pages in the local languages, **we expect *+local* and *+local+standard* to achieve better performance** than the *question-only* baseline and *+standard*. To isolate retrieval deficiencies attributable to cross-lingual retrieval rather than context information, we additionally assess (1e) *+local (translated)* on *yue-cmn* items: providing translated lead section (*yue*→*cmn*) by using Google Cloud Translation.⁶ However, due to the lack of existing high-quality sentence-level Bavarian-to-German translators (further discussed in Appendix B), we leave it to future work to evaluate the *+local (translated)* context for *bar-deu*.

For ECLeKTic, we used their translations of questions and answers into Mandarin/German and similarly tested three context scenarios: (2a) *question-only*: without any additional context information; (2c) *+source*: providing context text from the source Wikipedia article in another language;

⁶<https://cloud.google.com/translate>

Model	Context	yue-cmn					bar-deu				
		ROUGE-L	chrF++	BERTScore	BS-diff	LLMaJ	ROUGE-L	chrF++	BERTScore	BS-diff	LLMaJ
Llama3.1-8B	question-only	13.96	7.81	17.94	/	7.82	11.38	14.29	11.98	/	15.17
	+standard	14.65	8.20	19.25	+1.31	9.52	15.87	18.99	17.73	+5.75	18.54
	+local	36.63	25.75	43.24	+25.30	69.39	26.76	27.40	23.73	+11.75	62.36
	+standard+local	39.28	27.64	43.14	+25.20	65.31	26.21	27.65	26.28	+14.30	61.24
	+local (translated)	47.54	34.10	45.94	+28.00	71.43	/	/	/	/	/
Llama3.1-70B	question-only	20.60	12.47	26.57	/	14.63	17.89	20.94	23.75	/	24.16
	+standard	20.31	13.12	21.12	-5.45	18.37	21.04	23.24	27.15	+3.40	25.28
	+local	42.20	31.48	54.23	+27.66	85.03	28.77	27.33	31.13	+7.38	75.84
	+standard+local	45.30	33.35	52.95	+26.38	82.65	28.16	26.57	30.02	+6.27	73.03
	+local (translated)	56.47	41.81	55.86	+29.29	82.99	/	/	/	/	/
Qwen2.5-7B	question-only	14.40	7.76	19.17	/	11.22	18.53	25.55	26.32	/	12.36
	+standard	17.28	9.24	18.43	-0.74	13.61	19.43	25.49	25.47	-0.85	11.80
	+local	49.83	37.16	51.93	+32.76	81.97	25.81	28.88	30.62	+4.30	60.11
	+standard+local	48.92	36.31	50.34	+31.17	78.23	24.58	28.31	30.53	+4.21	56.74
	+local (translated)	55.42	40.97	53.88	+34.71	80.27	/	/	/	/	/
Qwen2.5-72B	question-only	17.05	11.35	23.17	/	16.33	18.56	23.39	26.96	/	17.42
	+standard	19.58	12.73	23.83	+0.66	20.75	22.91	28.52	29.63	+2.67	20.22
	+local	52.49	38.79	55.85	+32.68	88.10	35.03	36.82	40.28	+13.32	81.46
	+standard+local	53.54	40.42	54.94	+31.77	85.03	37.13	40.90	44.65	+17.69	76.97
	+local (translated)	58.03	42.74	55.92	+32.75	86.39	/	/	/	/	/
gpt-oos-20b	question-only	17.32	12.14	24.45	/	15.65	14.42	20.86	21.89	/	22.47
	+standard	19.59	13.54	26.88	+2.43	21.09	18.01	23.85	24.98	+3.09	25.28
	+local	48.06	35.60	50.74	+26.29	88.44	26.33	29.30	31.04	+9.15	70.79
	+standard+local	47.15	35.07	48.29	+23.84	86.39	25.14	27.84	29.59	+7.70	68.54
	+local (translated)	50.69	37.15	50.56	+26.11	85.37	/	/	/	/	/
gpt-oos-120b	question-only	20.91	14.41	26.24	/	20.41	17.75	23.85	25.98	/	23.03
	+standard	23.58	16.18	27.98	+1.74	32.65	20.33	26.35	28.06	+2.08	26.97
	+local	48.34	35.56	50.13	+23.89	89.46	28.90	32.41	34.81	+8.83	73.60
	+standard+local	48.21	35.34	49.01	+22.77	86.39	28.69	33.44	35.03	+9.05	73.03
	+local (translated)	51.33	37.73	51.12	+24.88	86.39	/	/	/	/	/

Table 2: Lead section QA performance, where the key information required to answer the questions is contained in the *+local* context. **Bold** indicates the highest score and **Blue** highlights the winner between *+local* and *+standard*, and all other settings higher than the winner. BS-diff measures the difference in BERTScore compared to the *question-only* baseline for the same LLM and prompt scenario.

and (2e) *+source (translated)*: providing the translated context text (others \rightarrow cmn/deu). As ECLeK-Tic addresses Wikipedia pages absent in the target languages, it lacks counterparts to (1b) and (1d).

Due to concerns on prompt-sensitive knowledge (i.e., facts retrieved only under specific prompt templates, Yin et al. 2024), we experiment with three prompt variations in our QA experiments. However, no significant performance difference is observed, and, for consistency, we report results using one of the prompt templates (see Appendix C).

Metrics We use standard evaluation metrics for Natural Language Generation (NLG) tasks to measure the similarity between the predicted and reference answers at the lexical, character, and semantic levels. ROUGE-L (Lin, 2004) evaluates lexical overlap of the longest common subsequence, chrF++ (Popović, 2015) captures character-level similarity that is robust across languages and morphological variants, and BERTScore (Zhang et al., 2020) measures token-level semantic similarity using contextualized embeddings from pretrained language models. We report F1 for ROUGE-L and BERTScore.

To complement traditional NLG evaluation metrics, we also employ LLM-as-a-judge (LLMaJ, Chiang and Lee 2023; Zheng et al. 2023). The judge model (gpt-oss-20b) takes the question, the reference answer, and the generated answer as input

and determines whether the generated answer is correct. We adapt the evaluation prompt from Goldman et al. (2025), which is provided in Appendix C. To verify the reliability and consistency of LLMaJ as an evaluation method, we manually inspected 50 QA instances evaluated by LLMaJ and found only one misjudgment (see Appendix D).

Lead section results Table 2 presents our main results on lead section QA. Adding a standard lead section (*+standard*) as context only shows marginal gain (e.g., +3.40 on bar-deu Llama3.1-70B) or even sometimes hurts performance (e.g., -5.45 on yue-cmn Llama3.1-70B), where *+local* substantially improves QA performance over the baseline (e.g., +27.66 on yue-cmn and +7.38 on bar-deu Llama3.1-70B). This agrees with our hypothesis that providing relevant context in the local language enhances the quality of LLM-generated answers. *+standard+local* generally achieves similar performance to *+local*, without further gains.

The improvements on yue-cmn are consistently larger than those observed on bar-deu. We hypothesize this might be due to LLMs’ poorer ability to extract relevant information from Bavarian, a lower-resource language than Cantonese. Table 3 illustrates an example of the generated answers by our best-performing model Qwen2.5-72B on the bar-deu example in Table 1. *+question-only* and *+standard* answers contain misinformation, in-

Question: Woher kommt der Name des Tanzes Hiataamadl? <i>Where does the name of the dance Hiataamadl come from?</i>
Reference Answer: Der Name des Tanzes kommt vom Tanzlied "Koa Hiataamadl mog i net". <i>The name of the dance comes from the dance tune "Koa Hiataamadl mog i net".</i>
question-only answer: Der Name des Tanzes Hiataamadl stammt aus der Mundart des Wiener Umlands und bedeutet so viel wie „hier tanzt man“ . <i>The name of the dance Hiataamadl comes from the dialect of the Viennese surroundings and means something like "here we dance".</i>
+standard answer: Der Name "Hiataamadl" stammt aus dem Bayerischen und setzt sich zusammen aus "Hiat" (was "Hügel" oder "Anhöhe" bedeutet) und "Madl" (was "Mädchen" bedeutet). Der Tanz wird also als "Tanz der Hügelmädchen" interpretiert, was auf die alpenländische Herkunft und die traditionelle Aufführung durch junge Frauen hinweist. <i>The name "Hiataamadl" comes from Bavarian and is composed of "Hiat" (meaning "hill" or "elevation") and "Madl" (meaning "girl"). The dance is therefore interpreted as "Dance of the Hill Girls," which refers to its Alpine origins and the traditional performance by young women.</i>
+local answer: Der Name des Tanzes Hiataamadl kommt vom Lied "Koa Hiataamadl mog i net". <i>The name of the dance Hiataamadl comes from the song "Koa Hiataamadl mog i net".</i>
+standard+local answer: Der Name des Tanzes Hiataamadl kommt vom gleichnamigen Lied "Koa Hiataamadl mog i net". <i>The name of the dance Hiataamadl comes from the song of the same name "Koa Hiataamadl mog i net".</i>

Table 3: Sample Qwen2.5-72B answers. *+local* and *+standard+local* answers are alike the reference, whereas *question-only* and *+standard* answers contain misinformation highlighted in purple.

cluding falsely claiming that the dance comes from "Viennese surroundings" and misinterpreting "Hiata" as hill, which means shepherd in Bavarian (More examples can be found in Appendix D).

Moreover, when local-to-standard translation is available, for $y_{ue} \rightarrow cmn$, *+local (translated)* tends to further improve scores on ROUGE-L, chrF++, and BERTScore, achieving the highest score. However, when measured by the accuracy of LLM-as-a-judge, *+local* achieves the highest performance, except Llama3.1-8B, where *+local (translated)* performs the best. This discrepancy may stem from the different nature of the metrics: LLMaJ directly assesses answer correctness, whereas the overlap-based metrics measure similarity with the reference text. These two evaluation perspectives are complementary, and we leave it to future work to systematically align these metrics with human judgments. Moreover, the logographic nature of Chinese characters (simplified for Mandarin and traditional for Cantonese) may lead to inconsistent evaluations between the two language groups at the character and token levels (ROUGE-L and chrF++). Therefore, we use the semantically based BERTScore for the follow-up results and analyses and will present other metrics in Appendix E.

Document-level results Table 4 presents results on document-level QA annotations to evaluate LLMs' capability in answering out-of-context questions. Adding additional context, such as *+standard*, *+local*, or *+local (translated)*, slightly improves performance over the *question-only* baseline. This minimal improvement is expected, as the added contexts lack the crucial information

Model	Context	yue-cmn	bar-deu
Llama3.1-8B	question-only	22.42	11.61
	+standard	+1.55	+1.18
	+local	-0.24	-0.40
	+standard+local	+5.06	+7.99
	+local (translated)	+1.06	/
Llama3.1-70B	question-only	24.48	22.39
	+standard	-0.49	+8.61
	+local	+0.42	+1.63
	+standard+local	+0.68	+5.59
	+local (translated)	+1.62	/
Qwen2.5-7B	question-only	23.09	27.39
	+standard	+0.80	+2.43
	+local	-1.08	-5.17
	+standard+local	+1.01	-6.12
	+local (translated)	-4.36	/
Qwen2.5-72B	question-only	25.83	25.36
	+standard	+2.00	+4.02
	+local	+3.25	+7.54
	+standard+local	+1.55	+6.64
	+local (translated)	+0.19	/
gpt-oss-20b	question-only	22.44	22.59
	+standard	-1.13	+3.95
	+local	+0.72	+6.32
	+standard+local	+1.24	+3.09
	+local (translated)	+3.59	/
gpt-oss-120b	question-only	23.23	23.38
	+standard	+3.05	+6.15
	+local	+1.72	+5.75
	+standard+local	+3.95	+7.51
	+local (translated)	+1.68	/

Table 4: Document-level BERTScore difference from *question-only* baseline; **bold** indicates best.

needed to answer the questions. Nevertheless, even context without answer-relevant information may provide limited benefits by helping the LLM better identify the topic or domain of the question. Overall, our results show that answer-relevant context is key to LLM QA performance.

ECLeKTic cross-lingual results We additionally report ECLeKTic's cross-lingual retrieval results in Figure 1. A clear and consistent pattern emerges: performance improves noticeably when the additional context contains the key information, even when the context language differs from the prompt language. The gains are even larger when the context is translated into the prompt language, consistent with the trend observed in the results of WiLOVA-QA's lead section. Because languages in the ECLeKTic dataset are relatively high-resource, performance on WiLOVA-QA (Table 2) is lower than on ECLeKTic (Figure 1), indicating that our low-resource varieties, Cantonese and Bavarian, pose greater challenges for LLMs.

5. Article Region and Topic

This section provides a fine-grained analysis of WiLOVA-QA's **geographical and topic diversity**. We conduct a stratified evaluation using these la-

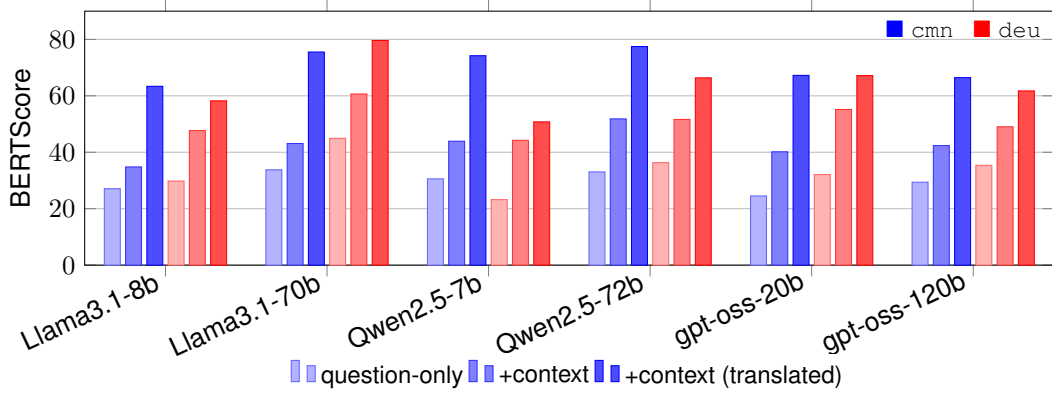


Figure 1: QA performance measured by BERTScore on the ECLeKTic dataset. Results compare context types of *question-only*, *+context*, and *+context (translated)* across models.

bels and investigate whether QA performance is more challenging for specific regions or topics.

Topic classification To characterize the topical coverage of our dataset, we assign a topic label to each selected article. Two leaders of the project (one fluent in Mandarin and one fluent in German) independently annotated the relevant language partition, and after adjudication, 14 topic categories were found relevant: *history*, *sport*, *geography*, *entertainment-art*, *animals-plants*, *politics-government*, *food*, *transportation*, *science-math-technology*, *linguistics*, *culture-customs*, *education*, *business*, and *other*.

Among the categories, *geography* (26 *yue-cmn* and 35 *bar-deu* articles) and *entertainment-art* (25 *yue-cmn* and 19 *bar-deu*) are the two most frequent in both languages. In our *bar-deu* data, *geography* includes different towns in Bavaria and Austria, but also outside of the Bavarian speaker regions, such as in Poland, the Faroe Islands, China, etc. *Entertainment-art* includes musicians, poets, movies, both in the local region and beyond. For *yue-cmn*, *geography* is also well balanced between street names and towns in the Cantonese-speaking regions, and location entities that are further away, such as in Spain, Australia, etc. *Entertainment* includes many TV shows and movies produced in Guangdong, Hong Kong, or Macau.

Region classification We also annotate each article for region relevance to understand the proportion of *yue* and *bar* articles that discuss concepts and cultures in the Cantonese- and Bavarian-speaking regions. Based on lead sections and Wikipedia categories, we annotate if the articles discuss information pertinent to the region where *local* variety is spoken, general to the *standard* language region, or *others*. *Local* includes regions such as Guangdong, Guangxi, Hong Kong, and Macau for *yue* and Bavaria, Austria, South Tyrol for

Topic/Region	Context	yue-cmn	bar-deu
<i>geography</i>	question-only	25.16	28.17
	+standard	22.99	31.08
	+local	57.02	37.23
	+standard+local	56.34	43.67
<i>entertainment-art</i>	question-only	24.00	22.79
	+standard	28.26	27.59
	+local	56.46	40.42
	+standard+local	57.18	44.07
<i>local region</i>	question-only	26.04	26.86
	+standard	26.26	31.65
	+local	59.55	43.46
	+standard+local	58.57	49.19
<i>all articles (avg)</i>	question-only	23.17	26.96
	+standard	23.83	29.63
	+local	55.85	40.28
	+standard+local	54.94	44.65

Table 5: Lead section BERTScore results of Qwen2.5-72B on top 2 topics and *local* articles.

bar. The *standard* region includes Greater China and Taiwan, as well as the rest of Germany, Italy, and Switzerland. Additionally, we define an *others* category to include articles pertinent to other specific regions, e.g., 16 *yue-cmn* articles discuss Japanese soccer players, as well as concepts that are region-independent, such as “triangle ruler,” “time,” or “Doppler effect.” Surprisingly, in both *yue-cmn* and *bar-deu* samples, less than half are *local* articles: 49 (35.8%) and 35 (37.2%). The majority are *other* articles that discuss concepts beyond the local or standard language region.

Results Table 5 presents Qwen2.5-72B (overall best performing model) results on the top 2 topic labels and *local*-region articles (full results on all topic and region labels are in Appendix F). In general, these stratified performances match the dataset’s average, indicating that *local*-relevant articles are **not** more challenging for LLMs than other articles. On *yue-cmn*, we observe that explicitly inserting lead sections from Cantonese achieves the best (or near best for *entertainment-*

art) performances. A recent evaluation benchmark, HKCanto-Eval (Cheng et al., 2025), also shows that Qwen2-72B (Yang et al., 2024) outperforms most models on academic, professional, and cultural QA for Cantonese. Overall, our WiLOVA-QA data provides a good testbed for low-resource evaluation scenarios of information asymmetry within and beyond regional knowledge.

6. Conclusion

This paper introduces WiLOVA-QA, a Wikipedia-based question-answering dataset that addresses information asymmetry across language varieties. Through experiments on Cantonese–Mandarin and Bavarian–German, we reveal that current LLMs often fail to recall local knowledge in closed-book settings, but can retrieve it effectively once the relevant local context is provided—highlighting a clear information asymmetry between high- and low-resource varieties. The findings suggest that knowledge gaps in LLMs parallel those observed on the web, where regional content remains underrepresented. Beyond QA, WiLOVA-QA offers a foundation for studying cross-variety understanding, cultural bias, and regional knowledge access in multilingual LLMs.

Acknowledgements

We thank Miriam Winkler and Yuchen Mao for their contributions and invaluable feedback on WiLOVA-QA. This work is funded by the ERC Consolidator Grant DIALECT 101043235.

7. References

- Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K. Arora, Yu Bai, Bowen Baker, Haiming Bao, Boaz Barak, Ally Bennett, Tyler Bertao, Nivedita Brett, Eugene Brevdo, Greg Brockman, Sebastien Bubeck, Che Chang, Kai Chen, Mark Chen, Enoch Cheung, Aidan Clark, Dan Cook, Marat Dukhan, Casey Dvorak, Kevin Fives, Vlad Fomenko, Timur Garipov, Kristian Georgiev, Mia Glaese, Tarun Gogineni, Adam Goucher, Lukas Gross, Katia Gil Guzman, John Hallman, Jackie Hehir, Johannes Heidecke, Alec Helyar, Haitang Hu, Romain Huet, Jacob Huh, Saachi Jain, Zach Johnson, Chris Koch, Irina Kofman, Dominik Kundel, Jason Kwon, Volodymyr Kyrylov, Elaine Ya Le, Guillaume Leclerc, James Park Lennon, Scott Lessans, Mario Lezcano-Casado, Yuanzhi Li, Zhuohan Li, Ji Lin, Jordan Liss, Lily, Liu, Jiancheng Liu, Kevin Lu, Chris Lu, Zoran Martinovic, Lindsay McCallum, Josh McGrath, Scott McKinney, Aidan McLaughlin, Song Mei, Steve Mostovoy, Tong Mu, Gideon Myles, Alexander Neitz, Alex Nichol, Jakub Pachocki, Alex Paino, Dana Palmie, Ashley Pantuliano, Giambattista Parascandolo, Jongsoo Park, Leher Pathak, Carolina Paz, Ludovic Peran, Dmitry Pimenov, Michelle Pokrass, Elizabeth Proehl, Huida Qiu, Gaby Raila, Filippo Raso, Hongyu Ren, Kimmy Richardson, David Robinson, Bob Rotsted, Hadi Salman, Suvansh Sanjeev, Max Schwarzer, D. Sculley, Harshit Sikchi, Kendal Simon, Karan Singhal, Yang Song, Dane Stuckey, Zhiqing Sun, Philippe Tillet, Sam Toizer, Foivos Tsimpourlas, Nikhil Vyas, Eric Wallace, Xin Wang, Miles Wang, Olivia Watkins, Kevin Weil, Amy Wendling, Kevin Whinnery, Cedric Whitney, Hannah Wong, Lin Yang, Yu Yang, Michihiro Yasunaga, Kristen Ying, Wojciech Zaremba, Wenting Zhan, Cyril Zhang, Brian Zhang, Eddie Zhang, and Shengjia Zhao. 2025. [gpt-oss-120b & gpt-oss-20b model card](#).
- Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pittler, Timothy Lillicrap, Angeliki Lazaridou, Orhan Firat, James Molloy, Michael Isard, Paul R. Barham, Tom Hennigan, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, Ryan Doherty, Eli Collins, Clemens Meyer, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, Jack Krawczyk, Cosmo Du, Ed Chi, Heng-Tze Cheng, Eric Ni, Purvi Shah, Patrick Kane, Betty Chan, Manaal Faruqui, Aliaksei Severyn, Hanzhao Lin, YaGuang Li, Yong Cheng, Abe Ittycheriah, Mahdis Mahdieh, Mia Chen, Pei Sun, Dustin Tran, Sumit Bagri, Balaji Lakshminarayanan, Jeremiah Liu, Andras Orban, Fabian Gura, Hao Zhou, Xinying Song, Aurelien Boffy, Harish Ganapathy, Steven Zheng, HyunJeong Choe, goston Weisz, Tao Zhu, Yifeng Lu, Siddharth Gopal, Jarrod Kahn, Maciej Kula, Jeff Pitman, Rushin Shah, Emanuel Taropa, Majd Al Merey, Martin Baeuml, Zhifeng Chen, Laurent El Shafey, Yujing Zhang, Olcan Sercinoglu, George Tucker, Enrique Piqueras, Maxim Krikun, Iain Barr, Nikolay Savinov, Ivo Danihelka, Becca Roelofs, Anaıs White, Anders Andreassen, Tamara von Glehn, Lakshman Yagati, Mehran Kazemi, Lucas Gonzalez, Misha Khalman, Jakub Sygnowski, Alexandre Frechette, Charlotte Smith, Laura Culp, Lev Proleev, Yi Luan, Xi Chen, James Lottes, Nathan Schucher, Federico Lebron, Alban Rustemi, Natalie Clay, Phil Crone, Tomas Kocisky, Jeffrey Zhao, Bartek Perz, Dian Yu, Heidi Howard,

Adam Bloniarz, Jack W. Rae, Han Lu, Laurent Sifre, Marcello Maggioni, Fred Alcober, Dan Garrette, Megan Barnes, Shantanu Thakoor, Jacob Austin, Gabriel Barth-Marón, William Wong, Rishabh Joshi, Rahma Chaabouni, Deeni Fatiha, Arun Ahuja, Gaurav Singh Tomar, Evan Senter, Martin Chadwick, Ilya Kornakov, Nithya Attaluri, Iñaki Iturrate, Ruibo Liu, Yunxuan Li, Sarah Cogan, Jeremy Chen, Chao Jia, Chenjie Gu, Qiao Zhang, Jordan Grimstad, Ale Jakse Hartman, Xavier Garcia, Thanumalayan Sankaranarayana Pillai, Jacob Devlin, Michael Laskin, Diego de Las Casas, Dasha Valter, Connie Tao, Lorenzo Blanco, Adrià Puigdomènech Badia, David Reitter, Mianna Chen, Jenny Brennan, Clara Rivera, Sergey Brin, Shariq Iqbal, Gabriela Surita, Jane Labanowski, Abhi Rao, Stephanie Winkler, Emilio Parisotto, Yiming Gu, Kate Olszewska, Ravi Addanki, Antoine Miech, Annie Louis, Denis Teplyashin, Geoff Brown, Elliot Catt, Jan Balaguer, Jackie Xiang, Pidong Wang, Zoe Ashwood, Anton Briukhov, Albert Webson, Sanjay Ganapathy, Smit Sanghavi, Ajay Kannan, Ming-Wei Chang, Axel Stjerngren, Josip Djolonga, Yuting Sun, Ankur Bapna, Matthew Aitchison, Pedram Pejman, Henryk Michalewski, Tianhe Yu, Cindy Wang, Juliette Love, Junwhan Ahn, Dawn Bloxwich, Kehang Han, Peter Humphreys, Thibault Sellam, James Bradbury, Varun Godbole, Sina Samangooei, Bogdan Damoc, Alex Kaskasoli, Sébastien M. R. Arnold, Vijay Vasudevan, Shubham Agrawal, Jason Riesa, Dmitry Lepikhin, Richard Tanburn, Srivatsan Srinivasan, Hyeontaek Lim, Sarah Hodgkinson, Pranav Shyam, Johan Ferret, Steven Hand, Ankush Garg, Tom Le Paine, Jian Li, Yujia Li, Minh Giang, Alexander Neitz, Zaheer Abbas, Sarah York, Machel Reid, Elizabeth Cole, Aakanksha Chowdhery, Dipanjan Das, Dominika Rogozińska, Vitaliy Nikolaev, Pablo Sprechmann, Zachary Nado, Lukas Zilka, Flavien Prost, Luheng He, Marianne Monteiro, Gaurav Mishra, Chris Welty, Josh Newlan, Dawei Jia, Miltiadis Allamanis, Clara Huiyi Hu, Raoul de Liedekerke, Justin Gilmer, Carl Saroufim, Shruti Rijhwani, Shaobo Hou, Disha Shrivastava, Anirudh Baddepudi, Alex Goldin, Adnan Ozturel, Albin Cassirer, Yunhan Xu, Daniel Sohn, Devendra Sachan, Reinald Kim Amplayo, Craig Swanson, Dessie Petrova, Shashi Narayan, Arthur Guez, Siddhartha Brahma, Jessica Landon, Miteyan Patel, Ruizhe Zhao, Kevin Vilella, Luyu Wang, Wenhao Jia, Matthew Rahtz, Mai Giménez, Legg Yeung, James Keeling, Petko Georgiev, Diana Mincu, Boxi Wu, Salem Haykal, Rachel Saputro, Kiran Vodrahalli, James Qin, Zeynep Cankara, Abhanshu Sharma, Nick Fernando, Will Hawkins, Behnam Neyshabur, Solomon

Kim, Adrian Hutter, Priyanka Agrawal, Alex Castro-Ros, George van den Driessche, Tao Wang, Fan Yang, Shuo yiin Chang, Paul Komarek, Ross McIlroy, Mario Lučić, Guodong Zhang, Wael Farhan, Michael Sharman, Paul Natsev, Paul Michel, Yamini Bansal, Siyuan Qiao, Kris Cao, Siamak Shakeri, Christina Butterfield, Justin Chung, Paul Kishan Rubenstein, Shivani Agrawal, Arthur Mensch, Kedar Soparkar, Karel Lenc, Timothy Chung, Aedan Pope, Loren Maggioro, Jackie Kay, Priya Jhakra, Shibo Wang, Joshua Maynez, Mary Phuong, Taylor Tobin, Andrea Tacchetti, Maja Trebacz, Kevin Robinson, Yash Katariya, Sebastian Riedel, Paige Bailey, Kefan Xiao, Nimesh Ghelani, Lora Aroyo, Ambrose Slone, Neil Houlsby, Xuehan Xiong, Zhen Yang, Elena Gribovskaya, Jonas Adler, Matteo Wirth, Lisa Lee, Music Li, Thais Kagohara, Jay Pavagadhi, Sophie Bridgers, Anna Bortsova, Sanjay Ghemawat, Zafarali Ahmed, Tianqi Liu, Richard Powell, Vijay Bolina, Mariko Iinuma, Polina Zablotskaia, James Besley, Da-Woon Chung, Timothy Dozat, Ramona Comanescu, Xiance Si, Jeremy Greer, Guolong Su, Martin Polacek, Raphaël Lopez Kaufman, Simon Tokumine, Hexiang Hu, Elena Buchatskaya, Yingjie Miao, Mohamed Elhawaty, Aditya Siddhant, Nenad Tomasev, Jinwei Xing, Christina Greer, Helen Miller, Shereen Ashraf, Aurko Roy, Zizhao Zhang, Ada Ma, Angelos Filos, Milos Besta, Rory Blevins, Ted Klimentko, Chih-Kuan Yeh, Soravit Changpinyo, Jiaqi Mu, Oscar Chang, Mantas Pajarskas, Carrie Muir, Vered Cohen, Charline Le Lan, Krishna Haridasan, Amit Marathe, Steven Hansen, Sholto Douglas, Rajkumar Samuel, Mingqiu Wang, Sophia Austin, Chang Lan, Jiepu Jiang, Justin Chiu, Jaime Alonso Lorenzo, Lars Lowe Sjösund, Sébastien Cevey, Zach Gleicher, Thi Avrahami, Anudhyan Boral, Hansa Srinivasan, Vittorio Selo, Rhys May, Konstantinos Aisopos, Léonard Hussenot, Livio Baldini Soares, Kate Baumli, Michael B. Chang, Adrià Recasens, Ben Caine, Alexander Pritzel, Filip Pavetic, Fabio Pardo, Anita Gergely, Justin Frye, Vinay Ramasesh, Dan Horgan, Kartikeya Badola, Nora Kassner, Subhrajit Roy, Ethan Dyer, Víctor Campos Campos, Alex Tomala, Yunhao Tang, Dalia El Badawy, Elspeth White, Basil Mustafa, Oran Lang, Abhishek Jindal, Sharad Vikram, Zhitao Gong, Sergi Caelles, Ross Hemsley, Gregory Thornton, Fangxiaoyu Feng, Wojciech Stokowiec, Ce Zheng, Phoebe Thacker, Çağlar Ünlü, Zhishuai Zhang, Mohammad Saleh, James Svensson, Max Bileschi, Piyush Patil, Ankesh Anand, Roman Ring, Katerina Tsihlias, Arpi Vezer, Marco Selvi, Toby Shevlane, Mikel Rodriguez, Tom Kwiatkowski, Samira Daruki, Keran Rong, Allan Dafoe, Nicholas

FitzGerald, Keren Gu-Lemberg, Mina Khan, Lisa Anne Hendricks, Marie Pellat, Vladimir Feinberg, James Cobon-Kerr, Tara Sainath, Maribeth Rauh, Sayed Hadi Hashemi, Richard Ives, Yana Hasson, Eric Noland, Yuan Cao, Nathan Byrd, Le Hou, Qingze Wang, Thibault Sottiaux, Michela Paganini, Jean-Baptiste Lespiau, Alexandre Moufarek, Samer Hassan, Kaushik Shivakumar, Joost van Amersfoort, Amol Mandhane, Pratik Joshi, Anirudh Goyal, Matthew Tung, Andrew Brock, Hannah Sheahan, Vedant Misra, Cheng Li, Nemanja Rakićević, Mostafa Dehghani, Fangyu Liu, Sid Mittal, Junhyuk Oh, Seb Noury, Eren Sezener, Fantine Huot, Matthew Lamm, Nicola De Cao, Charlie Chen, Sidharth Mudgal, Romina Stella, Kevin Brooks, Gautam Vasudevan, Chenxi Liu, Mainak Chain, Nivedita Melinkeri, Aaron Cohen, Venus Wang, Kristie Seymore, Sergey Zubkov, Rahul Goel, Summer Yue, Sai Krishnakumaran, Brian Albert, Nate Hurley, Motoki Sano, Anhad Mohananey, Jonah Joughin, Egor Filonov, Tomasz Kępa, Yomna Eldawy, Jiawern Lim, Rahul Rishi, Shirin Badiezadegan, Taylor Bos, Jerry Chang, Sanil Jain, Sri Gayatri Sundara Padmanabhan, Subha Puttagunta, Kalpesh Krishna, Leslie Baker, Norbert Kalb, Vamsi Bedapudi, Adam Kurzrok, Shuntong Lei, Anthony Yu, Oren Litvin, Xiang Zhou, Zhichun Wu, Sam Sobell, Andrea Siciliano, Alan Papir, Robby Neale, Jonas Bragnolo, Tej Toor, Tina Chen, Valentin Anklin, Feiran Wang, Richie Feng, Milad Gholami, Kevin Ling, Lijuan Liu, Jules Walter, Hamid Moghadam, Arun Kishore, Jakub Adamek, Tyler Mercado, Jonathan Mallinson, Siddhinita Wandekar, Stephen Cagle, Eran Ofek, Guillermo Garrido, Clemens Lombriser, Maksim Mukha, Botu Sun, Hafeezul Rahman Mohammad, Josip Matak, Yadi Qian, Vikas Peswani, Pawel Janus, Quan Yuan, Leif Schelin, Oana David, Ankur Garg, Yifan He, Oleksii Duzhyi, Anton Älgmyr, Timothée Lottaz, Qi Li, Vikas Yadav, Luyao Xu, Alex Chinien, Rakesh Shivanna, Aleksandr Chuklin, Josie Li, Carrie Spadine, Travis Wolfe, Kareem Mohamed, Subhabrata Das, Zihang Dai, Kyle He, Daniel von Dincklage, Shyam Upadhyay, Akanksha Maurya, Luyan Chi, Sebastian Krause, Khalid Salama, Pam G Rabinovitch, Pavan Kumar Reddy M, Aarush Selvan, Mikhail Dektiarev, Golnaz Ghiasi, Erdem Guven, Himanshu Gupta, Boyi Liu, Deepak Sharma, Idan Heimlich Shtacher, Shachi Paul, Oscar Akerlund, François-Xavier Aubet, Terry Huang, Chen Zhu, Eric Zhu, Elico Teixeira, Matthew Fritze, Francesco Bertolini, Liana-Eleonora Marinescu, Martin Bölle, Dominik Paulus, Khyatti Gupta, Tejasi Latkar, Max Chang, Jason Sanders, Roopa Wilson, Xuwei Wu, Yi-Xuan

Tan, Lam Nguyen Thiet, Tulsee Doshi, Sid Lall, Swaroop Mishra, Wanming Chen, Thang Luong, Seth Benjamin, Jasmine Lee, Ewa Andrejczuk, Dominik Rabiej, Vipul Ranjan, Krzysztof Styrz, Pengcheng Yin, Jon Simon, Malcolm Rose Harriott, Mudit Bansal, Alexei Robsky, Geoff Bacon, David Greene, Daniil Mirylenka, Chen Zhou, Obaid Sarvana, Abhimanyu Goyal, Samuel Andermatt, Patrick Siegler, Ben Horn, Assaf Israel, Francesco Pongetti, Chih-Wei "Louis" Chen, Marco Selvatici, Pedro Silva, Kathie Wang, Jackson Tolins, Kelvin Guu, Roey Yogev, Xiaochen Cai, Alessandro Agostini, Maulik Shah, Hung Nguyen, Noah Ó Donnaille, Sébastien Pereira, Linda Friso, Adam Stambler, Adam Kurzrok, Chenkai Kuang, Yan Romanikhin, Mark Geller, ZJ Yan, Kane Jang, Cheng-Chun Lee, Wojciech Fica, Eric Malmi, Qijun Tan, Dan Banica, Daniel Balle, Ryan Pham, Yanping Huang, Diana Avram, Hongzhi Shi, Jasjot Singh, Chris Hidey, Niharika Ahuja, Pranab Saxena, Dan Doolley, Srividya Pranavi Potharaju, Eileen O'Neill, Anand Gokulchandran, Ryan Foley, Kai Zhao, Mike Dusenberry, Yuan Liu, Pulkit Mehta, Ragha Kotikalapudi, Chalence Safranek-Shrader, Andrew Goodman, Joshua Kessinger, Eran Globen, Prateek Kolhar, Chris Gorgolewski, Ali Ibrahim, Yang Song, Ali Eichenbaum, Thomas Brovelli, Sahitya Potluri, Preethi Lahoti, Cip Baetu, Ali Ghorbani, Charles Chen, Andy Crawford, Shalini Pal, Mukund Sridhar, Petru Gurita, Asier Mujika, Igor Petrovski, Pierre-Louis Cedoz, Chenmei Li, Shiyuan Chen, Niccolò Dal Santo, Sidharth Goyal, Jitesh Punjabi, Karthik Kappaganthu, Chester Kwak, Pallavi LV, Sarmishta Velury, Himadri Choudhury, Jamie Hall, Premal Shah, Ricardo Figueira, Matt Thomas, Minjie Lu, Ting Zhou, Chintu Kumar, Thomas Jurdi, Sharat Chikkerur, Yenai Ma, Adams Yu, Soo Kwak, Victor Ähdel, Sujeewan Rajayogam, Travis Choma, Fei Liu, Aditya Barua, Colin Ji, Ji Ho Park, Vincent Hellendoorn, Alex Bailey, Taylan Bilal, Huanjie Zhou, Mehrdad Khatir, Charles Sutton, Wojciech Rzadkowski, Fiona Macintosh, Roopali Vij, Konstantin Shagin, Paul Medina, Chen Liang, Jinjing Zhou, Pararth Shah, Yingying Bi, Attila Dankovics, Shipra Banga, Sabine Lehmann, Marissa Bredesen, Zifan Lin, John Eric Hoffmann, Jonathan Lai, Raynald Chung, Kai Yang, Nihal Balani, Arthur Bražinskas, Andrei Sozanschi, Matthew Hayes, Héctor Fernández Alcalde, Peter Makarov, Will Chen, Antonio Stella, Liselotte Snijders, Michael Mandl, Ante Kärman, Paweł Nowak, Xinyi Wu, Alex Dyck, Krishnan Vaidyanathan, Raghavender R, Jessica Mallet, Mitch Rudominer, Eric Johnston, Sushil Mittal, Akhil Udathu, Janara Christensen, Vishal Verma, Zach Irving, Andreas Santucci, Gamaleldin El-

sayed, Elnaz Davoodi, Marin Georgiev, Ian Tenney, Nan Hua, Geoffrey Cideron, Edouard Laurent, Mahmoud Alnahlawi, Ionut Georgescu, Nan Wei, Ivy Zheng, Dylan Scandinaro, Heinrich Jiang, Jasper Snoek, Mukund Sundararajan, Xuezhong Wang, Zack Ontiveros, Itay Karo, Jeremy Cole, Vinu Rajashekhar, Lara Tume, Eyal Ben-David, Rishub Jain, Jonathan Uesato, Romina Datta, Oskar Bunyan, Shimu Wu, John Zhang, Piotr Stanczyk, Ye Zhang, David Steiner, Subhjit Naskar, Michael Azam, Matthew Johnson, Adam Paszke, Chung-Cheng Chiu, Jaume Sanchez Elias, Afroz Mohiuddin, Faizan Muhammad, Jin Miao, Andrew Lee, Nino Vieillard, Jane Park, Jiageng Zhang, Jeff Stanway, Drew Garmon, Abhijit Karmarkar, Zhe Dong, Jong Lee, Aviral Kumar, Luwei Zhou, Jonathan Evens, William Isaac, Geoffrey Irving, Edward Loper, Michael Fink, Isha Arkatkar, Nanxin Chen, Izhak Shafran, Ivan Petrychenko, Zhe Chen, Johnson Jia, Anselm Levskaya, Zhenkai Zhu, Peter Grabowski, Yu Mao, Alberto Magni, Kaisheng Yao, Javier Snider, Norman Casagrande, Evan Palmer, Paul Suganthan, Alfonso Castaño, Irene Giannoumis, Wooyeol Kim, Mikołaj Rybiński, Ashwin Sreevatsa, Jennifer Prendki, David Soergel, Adrian Goedeckemeyer, Willi Gierke, Mohsen Jafari, Meenu Gaba, Jeremy Wiesner, Diana Gage Wright, Yawen Wei, Harsha Vashisht, Yana Kulizhskaya, Jay Hoover, Maigo Le, Lu Li, Chimezie Iwuanyanwu, Lu Liu, Kevin Ramirez, Andrey Khorlin, Albert Cui, Tian LIN, Marcus Wu, Ricardo Aguilar, Keith Pallo, Abhishek Chakladar, Ginger Perng, Elena Allica Abellan, Mingyang Zhang, Ishita Dasgupta, Nate Kushman, Ivo Penchev, Alena Repina, Xihui Wu, Tom van der Weide, Priya Ponnapalli, Caroline Kaplan, Jiri Simsa, Shuangfeng Li, Olivier Dousse, Fan Yang, Jeff Piper, Nathan Ie, Rama Pasumarthi, Nathan Lintz, Anitha Vijayakumar, Daniel Andor, Pedro Valenzuela, Minnie Lui, Cosmin Paduraru, Daiyi Peng, Katherine Lee, Shuyuan Zhang, Somer Greene, Duc Dung Nguyen, Paula Kurylowicz, Cassidy Hardin, Lucas Dixon, Lili Janzer, Kiam Choo, Ziqiang Feng, Biao Zhang, Achintya Singhal, Dayou Du, Dan McKinnon, Natasha Antropova, Tolga Bolukbasi, Orgad Keller, David Reid, Daniel Finchelstein, Maria Abi Raad, Remi Crocker, Peter Hawkins, Robert Dadashi, Colin Gaffney, Ken Franko, Anna Bulanova, Rémi Leblond, Shirley Chung, Harry Askham, Luis C. Cobo, Kelvin Xu, Felix Fischer, Jun Xu, Christina Sorokin, Chris Alberti, Chu-Cheng Lin, Colin Evans, Alek Dimitriev, Hannah Forbes, Dylan Banarse, Zora Tung, Mark Omernick, Colton Bishop, Rachel Sterneck, Rohan Jain, Jiawei Xia, Ehsan Amid, Francesco Piccinno, Xingyu Wang, Praseem

Banzal, Daniel J. Mankowitz, Alex Polozov, Victoria Krakovna, Sasha Brown, MohammadHossein Bateni, Dennis Duan, Vlad Firoiu, Meghana Thotakuri, Tom Natan, Matthieu Geist, Ser tan Girgin, Hui Li, Jiayu Ye, Ofir Roval, Reiko Tojo, Michael Kwong, James Lee-Thorp, Christopher Yew, Danila Sinopalnikov, Sabela Ramos, John Mellor, Abhishek Sharma, Kathy Wu, David Miller, Nicolas Sonnerat, Denis Vnukov, Rory Greig, Jennifer Beattie, Emily Caveness, Libin Bai, Julian Eisenschlos, Alex Korchemniy, Tomy Tsai, Mimi Jasarevic, Weize Kong, Phuong Dao, Zeyu Zheng, Frederick Liu, Fan Yang, Rui Zhu, Tian Huey Teh, Jason Sanmiya, Evgeny Gladchenko, Nejc Trdin, Daniel Toyama, Evan Rosen, Sasan Tavakkol, Linting Xue, Chen Elkind, Oliver Woodman, John Carpenter, George Papamakarios, Rupert Kemp, Sushant Kafle, Tanya Grunina, Rishika Sinha, Alice Talbert, Diane Wu, Denese Owusu-Afryie, Cosmo Du, Chloe Thornton, Jordi Pont-Tuset, Pradyumna Narayana, Jing Li, Saaber Fatehi, John Wieting, Omar Ajmeri, Benigno Uria, Yeongil Ko, Laura Knight, Amélie Héliou, Ning Niu, Shane Gu, Chenxi Pang, Yeqing Li, Nir Levine, Ariel Stolovich, Rebeca Santamaria-Fernandez, Sonam Goenka, Wenny Yustalim, Robin Strudel, Ali Elqursh, Charlie Deck, Hyo Lee, Zonglin Li, Kyle Levin, Raphael Hoffmann, Dan Holtmann-Rice, Olivier Bachem, Sho Arora, Christy Koh, Soheil Hassas Yeganeh, Siim Põder, Mukarram Tariq, Yanhua Sun, Lucian Ionita, Mojtaba Seyedhosseini, Pouya Tafti, Zhiyu Liu, Anmol Gulati, Jasmine Liu, Xinyu Ye, Bart Chrzaszcz, Lily Wang, Nikhil Sethi, Tianrun Li, Ben Brown, Shreya Singh, Wei Fan, Aaron Parisi, Joe Stanton, Vinod Koverkathu, Christopher A. Choquette-Choo, Yunjie Li, TJ Lu, Abe Ittycheriah, Prakash Shroff, Mani Varadarajan, Sanaz Bahargam, Rob Willoughby, David Gaddy, Guillaume Desjardins, Marco Cornero, Brona Robenek, Bhavishya Mittal, Ben Albrecht, Ashish Shenoy, Fedor Moiseev, Henrik Jacobsson, Alireza Ghaffarkhah, Morgane Rivière, Alanna Walton, Clément Crepy, Alicia Parrish, Zongwei Zhou, Clement Farabet, Carey Radebaugh, Praveen Srinivasan, Claudia van der Salm, Andreas Fidjeland, Salvatore Scellato, Eri Latorre-Chimoto, Hanna Klimczak-Plucińska, David Bridson, Dario de Cesare, Tom Hudson, Piermaria Mendolicchio, Lexi Walker, Alex Morris, Matthew Mauger, Alexey Guseynov, Alison Reid, Seth Odoom, Lucia Loher, Victor Cotruta, Madhavi Yenugula, Dominik Grewe, Anastasia Petrushkina, Tom Duerig, Antonio Sanchez, Steve Yadlowsky, Amy Shen, Amir Globerson, Lynette Webb, Sahil Dua, Dong Li, Surya Bhatiraju, Dan Hurt, Haroon Qureshi, Ananth Agarwal, Tomer Shani, Matan Eyal, Anuj Khare,

- Shreyas Rammohan Belle, Lei Wang, Chetan Tekur, Mihir Sanjay Kale, Jinliang Wei, Ruoxin Sang, Brennan Saeta, Tyler Liechty, Yi Sun, Yao Zhao, Stephan Lee, Pandu Nayak, Doug Fritz, Manish Reddy Vuyyuru, John Aslanides, Nidhi Vyas, Martin Wicke, Xiao Ma, Evgenii Eltyshev, Nina Martin, Hardie Cate, James Manyika, Keyvan Amiri, Yelin Kim, Xi Xiong, Kai Kang, Florian Luisier, Nilesh Tripuraneni, David Madras, Mandy Guo, Austin Waters, Oliver Wang, Joshua Ainslie, Jason Baldridge, Han Zhang, Garima Pruthi, Jakob Bauer, Feng Yang, Riham Mansour, Jason Gelman, Yang Xu, George Polovets, Ji Liu, Honglong Cai, Warren Chen, Xiang-Hai Sheng, Emily Xue, Sherjil Ozair, Christof Angermueller, Xiaowei Li, Anoop Sinha, Weiren Wang, Julia Wiesinger, Emmanouil Koukoumidis, Yuan Tian, Anand Iyer, Madhu Gurusurthy, Mark Goldenson, Parashar Shah, MK Blake, Hongkun Yu, Anthony Urbanowicz, Jennimaria Palomaki, Chrisantha Fernando, Ken Durden, Harsh Mehta, Nikola Momchev, Elahe Rahimtoroghi, Maria Georgaki, Amit Raul, Sebastian Ruder, Morgan Redshaw, Jinhyuk Lee, Denny Zhou, Komal Jalan, Dinghua Li, Blake Hechtman, Parker Schuh, Milad Nasr, Kieran Milan, Vladimir Mikulik, Juliana Franco, Tim Green, Nam Nguyen, Joe Kelley, Aroma Mahendru, Andrea Hu, Joshua Howland, Ben Vargas, Jeffrey Hui, Kshitij Bansal, Vikram Rao, Rakesh Ghiya, Emma Wang, Ke Ye, Jean Michel Sarr, Melanie Moranski Preston, Madeleine Elish, Steve Li, Aakash Kaku, Jigar Gupta, Ice Pasupat, Da-Cheng Juan, Milan Someswar, Tejvi M., Xinyun Chen, Aida Amini, Alex Fabrikant, Eric Chu, Xuanyi Dong, Amruta Muthal, Senaka Buthpitiya, Sarthak Jauhari, Nan Hua, Urvashi Khandelwal, Ayal Hitron, Jie Ren, Larissa Rinaldi, Shahar Drath, Avigail Dabush, Nan-Jiang Jiang, Harshal Godhia, Uli Sachs, Anthony Chen, Yicheng Fan, Hagai Taitelbaum, Hila Noga, Zhuyun Dai, James Wang, Chen Liang, Jenny Hamer, Chun-Sung Ferng, Chenel Elkind, Aviel Atias, Paulina Lee, Vít Listík, Mathias Carlen, Jan van de Kerkhof, Marcin Pikus, Krunoslav Zaher, Paul Müller, Sasha Zykova, Richard Stefanec, Vitaly Gatsko, Christoph Hirschall, Ashwin Sethi, Xingyu Federico Xu, Chetan Ahuja, Beth Tsai, Anca Stefanoiu, Bo Feng, Keshav Dhandhanian, Manish Katyal, Akshay Gupta, Atharva Parulekar, Divya Pitta, Jing Zhao, Vivaan Bhatia, Yashodha Bhavnani, Omar Alhadlaq, Xiaolin Li, Peter Danenberg, Dennis Tu, Alex Pine, Vera Filippova, Abhipso Ghosh, Ben Limonchik, Bhargava Urala, Chaitanya Krishna Lanka, Derik Clive, Yi Sun, Edward Li, Hao Wu, Kevin Hongtongsak, Ianna Li, Kalind Thakkar, Kuanysh Omarov, Kushal Majmundar, Michael Alverson, Michael Kucharski, Mohak Patel, Mudit Jain, Maksim Zabelin, Paolo Pelagatti, Rohan Kohli, Saurabh Kumar, Joseph Kim, Swetha Sankar, Vineet Shah, Lakshmi Ramachandruni, Xiangkai Zeng, Ben Bariach, Laura Weidinger, Tu Vu, Alek Andreev, Antoine He, Kevin Hui, Sheleem Kashem, Amar Subramanya, Sissie Hsiao, Demis Hassabis, Koray Kavukcuoglu, Adam Sadovskiy, Quoc Le, Trevor Strohman, Yonghui Wu, Slav Petrov, Jeffrey Dean, and Oriol Vinyals. 2025. [Gemini: A family of highly capable multimodal models](#).
- Akari Asai, Jungo Kasai, Jonathan Clark, Kenton Lee, Eunsol Choi, and Hannaneh Hajishirzi. 2021. [XOR QA: Cross-lingual open-retrieval question answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 547–564, Online. Association for Computational Linguistics.
- Vahid Ashrafimoghari. 2023. [Detecting cross-lingual information gaps in wikipedia](#). In *Companion Proceedings of the ACM Web Conference 2023, WWW '23 Companion*, page 581–585, New York, NY, USA. Association for Computing Machinery.
- Patti Bao, Brent Hecht, Samuel Carton, Mahmood Quaderi, Michael Horn, and Darren Gergle. 2012. [Omnipedia: bridging the wikipedia language gap](#). In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '12*, page 1075–1084, New York, NY, USA. Association for Computing Machinery.
- Taryn Bipat, David W. McDonald, and Mark Zachry. 2018. [Do we all talk before we type? understanding collaboration in wikipedia language editions](#). In *Proceedings of the 14th International Symposium on Open Collaboration, OpenSym '18*, New York, NY, USA. Association for Computing Machinery.
- Michelle Brachman, Amina El-Ashry, Casey Dugan, and Werner Geyer. 2025. [Current and future use of large language models for knowledge work](#). *Proc. ACM Hum.-Comput. Interact.*, 9(7).
- Ewa S. Callahan and Susan C. Herring. 2011. [Cultural bias in wikipedia content on famous persons](#). *Journal of the American Society for Information Science and Technology*, 62(10):1899–1915.
- Tsz Chung Cheng, Chung Shing Cheng, Chaakming Lau, Eugene Lam, Wong Chun Yat, Hoi On Yu, and Cheuk Hei Chong. 2025. [HKCanto-eval: A benchmark for evaluating Cantonese language](#)

- understanding and cultural comprehension in LLMs. In *Proceedings of the 29th Conference on Computational Natural Language Learning*, pages 1–11, Vienna, Austria. Association for Computational Linguistics.
- Cheng-Han Chiang and Hung-yi Lee. 2023. [Can large language models be an alternative to human evaluations?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15607–15631, Toronto, Canada. Association for Computational Linguistics.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. [TyDi QA: A benchmark for information-seeking question answering in typologically diverse languages](#). *Transactions of the Association for Computational Linguistics*, 8:454–470.
- Brandon Duderstadt. 2025. [Wikivecs: A fully reproducible vectorization of multilingual Wikipedia](#). In *Proceedings of the 2nd Workshop on Advancing Natural Language Processing for Wikipedia (WikiNLP 2025)*, pages 1–9, Vienna, Austria. Association for Computational Linguistics.
- David M. Eberhard, Gary F. Simons, and Charles D. Fennig. 2025. [The ethnologue 200: What are the top 200 most spoken languages?](#) *Ethnologue Insights*. Online; accessed 2025-09-23.
- Jamie Elsey, Willem Slegers, and David Moss. 2025. [Estimating the usage and utility of llms in the us general public](#). Technical report, Rethink Priorities. Survey of 1,370 U.S. adults; accessed YYYY-MM-DD.
- Fahim Faisal, Sharlina Keshava, Md Mahfuz Ibn Alam, and Antonios Anastasopoulos. 2021. [SD-QA: Spoken dialectal question answering for the real world](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3296–3315, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Changjiang Gao, Hankun Lin, Xin Huang, Xue Han, Junlan Feng, Chao Deng, Jiajun Chen, and Shujian Huang. 2025. [Understanding LLMs’ cross-lingual context retrieval: How good it is and where it comes from](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 22797–22826, Suzhou, China. Association for Computational Linguistics.
- Omer Goldman, Uri Shaham, Dan Malkin, Sivan Eiger, Avinatan Hassidim, Yossi Matias, Joshua Maynez, Adi Mayrav Gilady, Jason Riesa, Shruti Rijhwani, Laura Rimell, Idan Szpektor, Reut Tsarfaty, and Matan Eyal. 2025. [Eclectic: a novel challenge set for evaluation of cross-lingual knowledge transfer](#).
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong,

Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vitor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuwei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina

Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelen, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho,

- Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. 2024. [The llama 3 herd of models](#).
- Julie May Groves. 2010. [Language or dialect, topolect or regiolect? A comparative study of language attitudes towards the status of Cantonese in Hong Kong](#). *Journal of Multilingual and Multicultural Development*, 31(6):531–551.
- Wan-hua Her and Udo Kruschwitz. 2024. [Investigating neural machine translation for low-resource languages: Using Bavarian as a case study](#). In *Proceedings of the 3rd Annual Meeting of the Special Interest Group on Under-resourced Languages @ LREC-COLING 2024*, pages 155–167, Torino, Italia. ELRA and ICCL.
- Michael Hoffmann, Jophin John, Stefan Schweter, Gokul Ramakrishnan, Hoi-Fong Mak, Alice Zhang, Dmitry Gaynullin, and Nicolay J. Hammer. 2025. [Llama-genba-10b: A trilingual large language model for german, english and bavarian](#).
- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020. [How can we know what language models know?](#) *Transactions of the Association for Computational Linguistics*, 8:423–438.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The state and fate of linguistic diversity and inclusion in the NLP world](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Nora Kassner and Hinrich Schütze. 2020. [Negated and misprimed probes for pretrained language models: Birds can talk, but cannot fly](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7811–7818, Online. Association for Computational Linguistics.
- Josef Kolbitsch and Hermann A Maurer. 2006. The transformation of the web: How emerging communities shape the information we consume. *J. Univers. Comput. Sci.*, 12(2):187–213.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Włodzimierz Lewoniewski, Krzysztof Węcel, and Witold Abramowicz. 2017. Analysis of references across wikipedia languages. In *Information and Software Technologies*, pages 561–573, Cham. Springer International Publishing.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Robert Litschko, Verena Blaschke, Diana Burkhardt, Barbara Plank, and Diego Frassinelli. 2025a. [Make every letter count: Building dialect variation dictionaries from monolingual corpora](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 14157–14174, Suzhou, China. Association for Computational Linguistics.
- Robert Litschko, Oliver Kraus, Verena Blaschke, and Barbara Plank. 2025b. [Cross-dialect information retrieval: Information access in low-resource and high-variance languages](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 10158–10171, Abu Dhabi, UAE. Association for Computational Linguistics.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822, Toronto, Canada. Association for Computational Linguistics.
- Junho Myung, Nayeon Lee, Yi Zhou, Jiho Jin, Rifki Putri, Dimosthenis Antypas, Hsuvas Borkakoty, Eunsu Kim, Carla Perez-Almendros, Abinew Ali Ayele, et al. 2024. Blend: A benchmark for llms on everyday knowledge in diverse cultures and languages. *Advances in Neural Information Processing Systems*, 37:78104–78146.
- Odunayo Ogundepo, Tajuddeen R. Gwadabe, Clara E. Rivera, Jonathan H. Clark, Sebastian

- Ruder, David Ifeoluwa Adelani, Bonaventure F. P. Dossou, Abdou Aziz Diop, Claytone Sikasote, Gilles Hacheme, Happy Buzaaba, Ignatius Ezeani, Rooweither Mabuya, Salomey Osei, Chris Emezue, Albert Njoroge Kahira, Shamsuddeen Hassan Muhammad, Akintunde Oladipo, Abraham Toluwase Owodunni, Atnafu Lambebo Tonja, Iyanuoluwa Shode, Akari Asai, Tunde Oluwaseyi Ajayi, Clemencia Siro, Steven Arthur, Mofetoluwa Adeyemi, Orevaoghene Ahia, Anuoluwapo Aremu, Oyinkansola Awosan, Chiamaka Chukwuneke, Bernard Opoku, Awokoya Ayodele, Verrah Otiende, Christine Mwase, Boyd Sinkala, Andre Niyongabo Rubungo, Daniel A. Ajisafe, Emeka Felix Onwuegbuzia, Habib Mbow, Emile Niyomutabazi, Eunice Mukonde, Falalu Ibrahim Lawan, Ibrahim Said Ahmad, Jesujoba O. Alabi, Martin Namukombo, Mbonu Chinedu, Mofya Phiri, Neo Putini, Ndumiso Mngoma, Priscilla A. Amouk, Ruqayya Nasir Iro, and Sonia Adhiambo. 2023. [AfriQA: Cross-lingual open-retrieval question answering for African languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14957–14972, Singapore. Association for Computational Linguistics.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. [How much knowledge can you pack into the parameters of a language model?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5418–5426, Online. Association for Computational Linguistics.
- Angelika Romanou, Negar Foroutan, Anna Sotnikova, Sree Harsha Nelaturu, Shivalika Singh, Rishabh Maheshwary, Micol Altomare, Zeming Chen, Mohamed A. Haggag, Snegha A, Alfonso Amayuelas, Azril Hafizi Amirudin, Danylo Boiko, Michael Chang, Jenny Chim, Gal Cohen, Aditya Kumar Dalmia, Abraham Diress, Sharad Duwal, Daniil Dzenhaliou, Daniel Fernando Eraso Florez, Fabian Farestam, Joseph Marvin Imperial, Shayekh Bin Islam, Perttu Isotalo, Maral Jabbarishiviari, Börje F. Karlsson, Eldar Khalilov, Christopher Klamm, Fajri Koto, Dominik Krzemiński, Gabriel Adriano de Melo, Syrielle Montariol, Yiyang Nan, Joel Niklaus, Jekaterina Novikova, Johan Samir Obando Ceron, Debjit Paul, Esther Ploeger, Jebish Purbey, Swati Rajwal, Selvan Sunitha Ravi, Sara Rydell, Roshan Santhosh, Drishti Sharma, Marjana Prifti Skenduli, Arshia Soltani Moakhar, Bardia soltani moakhar, Ayush Kumar Tarun, Azmine Touthik Wasi, Thenuka Ovin Weerasinghe, Serhan Yilmaz, Mike Zhang, Imanol Schlag, Marzieh Fadaee, Sara Hooker, and Antoine Bosselut. 2025. [INCLUDE: Evaluating multilingual language understanding with regional knowledge](#). In *The Thirteenth International Conference on Learning Representations*.
- David Romero, Chenyang Lyu, Haryo Akbarianto Wibowo, Teresa Lynn, Injy Hamed, Aditya Nanda Kishore, Aishik Mandal, Alina Dragonetti, Artem Abzaliev, Atnafu Lambebo Tonja, Bontu Fufa Balcha, Chenxi Whitehouse, Christian Salamea, Dan John Velasco, David Ifeoluwa Adelani, David Le Meur, Emilio Villacueva, Fajri Koto, Fauzan Farooqui, Federico Belcavello, Ganzorig Batnasan, Gisela Vallejo, Grainne Caulfield, Guido Ivetta, Haiyue Song, Henok Biadgign Ademtew, Hernán Maina, Holy Lovenia, Israel Abebe Azime, Jan Christian Blaise Cruz, Jay Gala, Jiahui Geng, Jesus-German Ortiz-Barajas, Jinheon Baek, Jocelyn Dunstan, Laura Alonso Alemany, Kumaranage Ravindu Yasas Nagasinghe, Luciana Benotti, Luis Fernando D’Haro, Marcelo Viridiano, Marcos Estecha-Garitagotia, Maria Camila Buitrago Cabrera, Mario Rodríguez-Cantelar, Mélanie Joutiteau, Mihail Mihaylov, Naome Etori, Mohamed Fazli Mohamed Imam, Muhammad Farid Adilazuarda, Munkhjargal Gochoo, Munkh-Erdene Otgonbold, Olivier Niyomugisha, Paula Mónica Silva, Pranjal Chitale, Raj Dabre, Rendi Chevi, Ruochen Zhang, Ryandito Diandaru, Samuel Cahyawijaya, Santiago Góngora, Soyeong Jeong, Sukannya Purkayastha, Tatsuki Kuribayashi, Teresa Clifford, Thanmay Jayakumar, Tiago Timponi Torrent, Toqeer Ehsan, Vladimir Araujo, Yova Kementchedjheva, Zara Burzo, Zheng Wei Lim, Zheng Xin Yong, Oana Ignat, Joan Nwatu, Rada Mihalcea, Tamar Solorio, and Alham Fikri Aji. 2025. [Cvqa: culturally-diverse multilingual visual question answering benchmark](#). In *Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS ’24*, Red Hook, NY, USA. Curran Associates Inc.

- Dwaipayan Roy, Sumit Bhatia, and Prateek Jain. 2020. [A topic-aligned multilingual corpus of Wikipedia articles for studying information asymmetry in low resource languages](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 2373–2380, Marseille, France. European Language Resources Association.
- Dwaipayan Roy, Sumit Bhatia, and Prateek Jain. 2022. Information asymmetry in wikipedia across different languages: A statistical analysis. *Journal of the Association for Information Science and Technology*, 73(3):347–361.
- Farhan Samir, Chan Young Park, Anjalie Field, Vered Shwartz, and Yulia Tsvetkov. 2024. [Locating information gaps and narrative inconsistencies across languages: A case study of LGBT people portrayals on Wikipedia](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6747–6762, Miami, Florida, USA. Association for Computational Linguistics.
- Nikhil Sharma, Kenton Murray, and Ziang Xiao. 2025. [Faux polyglot: A study on information disparity in multilingual large language models](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8090–8107, Albuquerque, New Mexico. Association for Computational Linguistics.
- Zining Wang, Yuxuan Zhang, Dongwook Yoon, Nicholas Vincent, Farhan Samir, and Vered Shwartz. 2025. Wikigap: Promoting epistemic equity by surfacing knowledge gaps between english wikipedia and other language editions. *arXiv preprint arXiv:2505.24195*.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. 2024. [Qwen2 technical report](#).
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025a. [Qwen2.5 technical report](#).
- Wenli Yang, Lilian Some, Michael Bain, and Byeong Kang. 2025b. [A comprehensive survey on integrating large language models with knowledge-based methods](#). *Knowledge-Based Systems*, 318:113503.
- Howard Yen, Tianyu Gao, Jinhyuk Lee, and Danqi Chen. 2023. [MoQA: Benchmarking multi-type open-domain question answering](#). In *Proceedings of the Third DialDoc Workshop on Document-grounded Dialogue and Conversational Question Answering*, pages 8–29, Toronto, Canada. Association for Computational Linguistics.
- Xunjian Yin, Xu Zhang, Jie Ruan, and Xiaojun Wan. 2024. [Benchmarking knowledge boundary for large language models: A different perspective on model evaluation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2270–2286, Bangkok, Thailand. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Xiang Zhang, Senyu Li, Bradley Hauer, Ning Shi, and Grzegorz Kondrak. 2023. [Don't trust ChatGPT when your question is not in English: A study of multilingual abilities and types of LLMs](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7915–7927, Singapore. Association for Computational Linguistics.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 46595–46623. Curran Associates, Inc.

8. Language Resource References

Omer Goldman, Uri Shaham, Dan Malkin, Sivan Eiger, Avinatan Hassidim, Yossi Matias, Joshua Maynez, Adi Mayrav Gilady, Jason Riesa, Shruti Rijhwani, Laura Rimell, Idan Szpektor, Reut Tsarfaty, and Matan Eyal. 2025. [Eclctic: a novel challenge set for evaluation of cross-lingual knowledge transfer](#).

A. WiLoVA-QA statistics

language	Bavarian	German	Cantonese	Mandarin
ISO 639-3	bar	deu	yue	cmn
#speakers (#rank) ⁷	13.7M (#92)	134.0M (#12)	85.7M (#25)	1.2B (#2)
#wiki-pages ⁸	27.2K	3,053.0K	146.3K	1,501.5K
#aligned-pages	24.6K		111.6K	
#local-heavy-pages	1,038		6,173	
#inspected-pages	497		229	
#qa-pages	94		137	
#lead section-qa	178		294	
#doc-qa	46		80	

Table 6: Statistics of the WiLoVA-QA dataset and included language varieties.

B. Challenges in Bavarian–German translation

Machine translation involving Bavarian remains largely underexplored. Prior work mainly relies on LLMs to translate from Bavarian to German (Litschko et al., 2025a; Hoffmann et al., 2025), and to the best of our knowledge, the only study training a dedicated MT system for Bavarian is Her and Kruschwitz (2024). In contrast, high-performing translation systems exist for Cantonese. To avoid confounding our question-answering results with translation errors, we conduct translation experiments only on Cantonese.

C. Prompt Templates

- **Prompt template: baseline**
Please answer this question: {question}
Please enclose your final answer in <Answer>...</Answer> tags. Please make sure to answer in English.
- **Prompt template: standard language summary**
Please answer this question: {question}
Here is some relevant information: {standard language summary}
Please enclose your final answer in <Answer>...</Answer> tags. Please make sure to answer in English.

- **Prompt template: dialect summary**
Please answer this question: {question}
Here is some relevant information: {dialect summary}
Please enclose your final answer in <Answer>...</Answer> tags. Please make sure to answer in English.
- **Prompt template: standard language summary + dialect summary**
Please answer this question: {question}
Here is some relevant information: {standard language summary}
Here is some relevant information: {dialect summary}
Please enclose your final answer in <Answer>...</Answer> tags. Please make sure to answer in English.
- **Prompt template of LLM-as-a-judge**
Your task is to evaluate whether a generated answer correctly answers the question, using the provided ground truth answer as reference.
Question: {question}
Ground Truth Answer: {ground truth}
Generated Answer: {generated answer}
Please output a single capitalized word (YES or NO) as evaluation result, without any additional explanation:
- YES: the generated answer is correct according to the ground truth answer.
- NO: the generated answer is incorrect according to the ground truth answer.

D. Error examples

- **LLMaJ Misjudgment**
index: 204
question: 肥冢一晃曾效力过哪些俱乐部? (*Which clubs has Kazuaki Koezuka played for?*)
answer: 大阪飞脚, 京都不死鸟 (*Gamba Osaka, Kyoto Sanga*)
ground truth answer: 大阪飞脚和京都不死鸟。 (*Gamba Osaka and Kyoto Sanga.*)
LLMaJ: NO
- **QA error even when provided with sufficient context (gpt-oss-120B)**
index: 170
question: 《五个小孩的大叔》的节目主持是谁? (*Who is the host of the program "5 Kids and a Bloke"?*)
prompt: 请回答以下问题: 《五个小孩的大叔》的节目主持是谁? 这里有一些相关的信息: 《五个小孩的大叔》(英语: 5 Kids And A Bloke, 粤文意思: 五个孩子的大叔)

叔) 是香港电视广播有限公司制作的真人骚节目, 全节目共四集, 在2016年10月2号到23号逢礼拜日晚20:00-20:30在翡翠台播, 并在myTV提供节目重温。节目主持是洪永城和电影《五个小孩的校长》里五个小主角: 何、傅舜盈、王诗雅、李咏珊、陈丽儿。请用中文回答, 并且将你的最终答案放在<Answer> 和</Answer> 标签之间。在标签内只保留最终答案内容, 不要包含任何多余解释或其他文字。(Please answer the following question: Who is the host of the program "5 Kids and a Bloke"? Here is some relevant information: "5 Kids and a Bloke" (English: 5 Kids And A Bloke; in Cantonese, meaning "Five Kids and an Uncle") is a reality show produced by Television Broadcasts Limited (TVB), consisting of four episodes. It was broadcast on Jade Channel every Sunday night from October 2 to October 23, 2016 (20:00-20:30) and made available for replay on myTV. The hosts are Tony Hung and the five child actors from the movie "Little Big Master": Ho Yuen-ying, Fu Shun-ying, Wong Sze-nga, Lee Wing-shan, and Chan Lai-ye. Please answer in Chinese and put your final answer between the <Answer> and </Answer> tags. Inside the tags, include only the final answer, without any explanations or extra text.)

answer: 洪永城(Tony Hung)

ground truth answer: 是洪永城和电影《五个小孩的校长》中的五个小主角: 何滢、傅舜盈、王诗雅、李咏珊、陈丽儿。(Tony Hung and the five child actors from the movie "Little Big Master": Ho Yuen-ying, Fu Shun-ying, Wong Sze-nga, Lee Wing-shan, and Chan Lai-ye.)

LLMaJ: NO

• **QA answer comparison across different contexts for the same model (Llama3.1-70B): Cantonese-Mandarin**

question: 进奏院广东话读音的粤拼是怎样的? (What is the Cantonese romanization, of the word "进奏院"?)

question-only: zin3 zau3 jyun2

answer(+standard): jin4 zhou4 jyun2

answer(+local): zeon3 zau3 jyun2

answer(+standard+local): zeon3 zau3 jyun2

answer(+local(translated)): zeon3 zau3 jyun2

ground truth answer: 进奏院广东话读音的粤拼是zeon3 zau3 jyun2。(The Cantonese romanization for "进奏院" is zeon3 zau3 jyun2.)

• **QA answer comparison across different contexts for the same model (Llama3.1-70B): Bavarian-German**

question: Wo wurde 1983 schon englischer Rap in der Schweiz gespielt? (Where was English rap already played in Switzerland in 1983?)

question-only: Basel (Basel)

answer(+standard): Basel (Basel)

answer(+local): Auf dem Radio-Sender DRS 3. (On the radio station DRS 3.)

answer(+standard+local): Auf Radio-Senda DRS 3 (On the radio station DRS 3.)

ground truth answer: 1983 wurde auf dem Radiosender DRS 3 schon englischer Rap gespielt. (In 1983, English rap was already played on the radio station DRS 3.)

• **QA answer comparison for the same context type (+local) across different models**

index: 15

question: 爱尔兰国防军的总司令是谁? 真正形式总司令职权的是谁? (Who is the Commander-in-Chief of the Irish Defence Forces? Who actually exercises the powers of the Commander-in-Chief?)

prompt: 请回答以下问题: 爱尔兰国防军的总司令是谁? 真正形式总司令职权的是谁? 这里有一些相关的信息: 愛爾蘭國防軍 (愛爾蘭蓋爾文: Óglaigh na hÉireann、英文: Irish Defence Forces) 係愛爾蘭共和國武裝力量, 由陸軍、海上部隊、空中部隊同後備國防軍四個部份組成。愛爾蘭國防軍總司令係愛爾蘭總統, 但係以總統名義行使總司令職權通常係國防部長, 而國防部長諮詢機構係國防委員會。愛爾蘭國防軍個愛爾蘭蓋爾文名, 同愛爾蘭共和軍愛爾蘭蓋爾文名完全一樣, 反映共和軍同國防軍之間上傳下承歷史關係。请用中文回答, 并且将你的最终答案放在<Answer> 和</Answer> 标签之间。在标签内只保留最终答案内容, 不要包含任何多余解释或其他文字。(Please answer the following question: Who is the Commander-in-Chief of the Irish Defence Forces? Who actually exercises the powers of the Commander-in-Chief? Here is some relevant information: The Irish Defence Forces (Irish: Óglaigh na hÉireann, English: Irish Defence Forces) are the armed forces of the Republic of Ireland, consisting of the Army, Naval Service, Air Corps, and Reserve Defence Forces. The Commander-in-Chief of the Irish Defence Forces is the President of Ireland, but the powers of the Commander-in-Chief are usually exercised in the President's

name by the Minister for Defence, whose advisory body is the Defence Council. The Irish name of the Defence Forces is identical to that of the Irish Republican Army, reflecting the historical continuity between the IRA and the Defence Forces. Please answer in Chinese and place your final answer between the <Answer> and </Answer> tags. Inside the tags, include only the final answer without any additional explanation or text.)

llama3.1-8b answer: 爱尔兰总统(*President of Ireland*)

llama3.1-70b answer: 愛爾蘭總統 (*President of Ireland*)

qwen2.5-7b answer: 愛爾蘭总统(*President of Ireland*)

qwen2.5-72b answer: 愛爾蘭總統 (*President of Ireland*)

gpt-oss-20b answer: 爱尔兰国防军的总司令是爱尔兰总统；真正行使总司令职权的是国防部长。(The *Commander-in-Chief of the Irish Defence Forces is the President of Ireland; the one who actually exercises the powers of the Commander-in-Chief is the Minister for Defence.*)

gpt-oss-120b answer: 爱尔兰国防军的名义总司令是爱尔兰总统，但实际行使总司令职权的通常是国防部长。(The *nominal Commander-in-Chief of the Irish Defence Forces is the President of Ireland, but the powers are usually exercised in practice by the Minister for Defence.*)

ground truth answer: 爱尔兰国防军的总司令是爱尔兰总统，但是以总统的名义行使总司令职权的通常是国防部长，而国防部长的谘询机构是国防委员会。(The *Commander-in-Chief of the Irish Defence Forces is the President of Ireland, but the powers of the Commander-in-Chief are usually exercised in the President's name by the Minister for Defence, whose advisory body is the Defence Council.*)

E. Complete results of all metrics

Model	Context	yue-cmn				bar-deu			
		rouge-l	chrF++	BERTScore	LLMaJ	rouge-l	chrF++	BERTScore	LLMaJ
Llama3.1-8B	question-only	15.66	7.02	22.42	7.50	10.49	12.86	11.61	8.70
	+standard	19.41	9.24	23.97	13.75	13.31	14.87	12.79	4.35
	+local	17.69	7.90	22.18	8.75	11.57	13.32	11.21	4.35
	+standard+local	23.16	10.16	27.48	12.50	19.77	21.11	19.60	6.52
	+local (translated)	17.96	8.33	23.48	8.75	/	/	/	/
Llama3.1-70B	question-only	18.85	8.05	24.48	12.50	15.34	17.92	22.39	15.22
	+standard	19.44	10.22	23.99	12.50	24.76	26.29	31.00	21.74
	+local	19.69	10.69	24.90	16.25	20.62	22.67	24.02	13.04
	+standard+local	21.53	10.39	25.16	16.25	24.19	25.87	27.98	19.57
	+local (translated)	22.08	10.77	26.10	16.25	/	/	/	/
Qwen2.5-7B	question-only	16.62	8.52	23.09	12.50	21.96	27.53	27.39	0.00
	+standard	20.93	11.01	23.89	12.50	24.68	31.15	29.82	6.52
	+local	19.66	8.94	22.01	11.25	17.32	24.45	22.22	2.17
	+standard+local	20.60	9.06	24.10	11.25	19.50	24.50	21.27	4.35
	+local (translated)	17.50	8.51	18.73	7.50	/	/	/	/
Qwen2.5-72B	question-only	19.99	9.63	25.83	17.50	18.80	24.58	25.36	10.87
	+standard	20.57	10.54	27.83	15.00	22.69	29.87	29.38	15.22
	+local	21.96	10.65	29.08	16.25	26.58	33.43	32.90	6.52
	+standard+local	22.13	11.11	27.38	17.50	26.15	33.43	32.00	15.22
	+local (translated)	20.17	10.59	26.02	18.75	/	/	/	/
gpt-oss-20b	question-only	15.52	8.27	22.44	12.50	15.84	20.92	22.59	8.70
	+standard	16.42	8.36	21.31	15.00	18.68	23.80	26.54	8.70
	+local	18.04	9.54	23.16	16.25	20.19	25.28	28.91	13.04
	+standard+local	17.80	10.17	23.68	18.75	18.30	23.70	25.68	13.04
	+local (translated)	19.49	10.40	26.03	15.00	/	/	/	/
gpt-oss-120b	question-only	15.94	9.58	23.23	15.00	16.47	19.74	23.38	13.04
	+standard	19.47	11.25	26.28	23.75	21.63	26.68	29.53	15.22
	+local	18.93	11.21	24.95	26.25	20.72	26.64	29.13	17.39
	+standard+local	22.58	12.94	27.18	28.75	24.96	30.93	30.89	13.04
	+local (translated)	20.61	11.60	24.91	25.00	/	/	/	/

Table 7: Document-level out-of-context QA performance as measured by all metrics. Note that all added contexts lack the information needed to answer the questions.

Model	Context	cmn				deu			
		rouge-l	chrF++	BERTScore	LLMaJ	rouge-l	chrF++	BERTScore	LLMaJ
Llama3.1-8b	question-only	17.67	12.15	27.08	10.51	9.97	16.42	29.81	13.81
	+context	34.27	31.38	34.80	57.36	41.87	44.13	47.69	62.46
	+context (translated)	63.83	59.88	63.37	72.37	62.07	64.03	58.22	79.88
Llama3.1-70b	question-only	25.49	19.29	33.79	21.62	23.44	30.08	44.94	27.33
	+context	42.03	38.45	43.09	72.37	49.83	52.67	60.65	76.88
	+context (translated)	74.79	71.39	75.52	86.79	76.12	79.15	79.61	86.79
Qwen2.5-7b	question-only	21.11	13.69	30.57	12.31	8.17	15.36	23.21	7.81
	+context	42.56	37.82	43.91	63.66	40.46	46.51	44.22	67.27
	+context (translated)	73.61	70.82	74.21	85.89	48.33	56.36	50.77	81.38
Qwen2.5-72b	question-only	23.63	18.22	33.04	19.22	14.21	23.02	36.30	16.82
	+context	48.57	44.33	51.82	72.67	46.65	54.14	51.62	80.48
	+context (translated)	77.09	73.57	77.47	88.59	63.71	70.49	66.37	90.69
gpt-oss-20b	question-only	15.37	13.28	24.52	15.62	10.75	18.31	32.09	15.02
	+context	39.84	36.24	40.16	76.88	49.92	57.16	55.16	78.68
	+context (translated)	66.91	63.63	67.24	86.79	63.11	69.20	67.17	88.29
gpt-oss-120b	question-only	18.99	16.23	29.38	22.52	14.23	22.69	35.38	20.72
	+context	41.54	38.44	42.36	79.88	44.95	53.53	49.02	82.58
	+context (translated)	65.90	62.31	66.44	87.39	58.83	66.05	61.74	88.89

Table 8: QA performance measured by all metrics on the ECLeKTic dataset. Results compare context types of *question-only*, *+context*, and *+context (translated)* across models.

F. Topic and region results

Topic	Context	yue-cmn				bar-deu			
		rouge-l	chrF++	BERTScore	LLMaJ	rouge-l	chrF++	BERTScore	LLMaJ
science-math-technology	question-only	19.59	11.15	22.33	22.58	10.76	19.80	13.01	33.33
	+standard	20.82	13.34	26.42	38.71	13.47	20.25	12.71	33.33
	+local	52.86	36.85	53.71	87.10	14.05	22.69	17.89	66.67
	+standard+local	52.74	39.43	48.60	90.32	17.28	27.27	21.12	66.67
geography	question-only	15.36	11.75	25.16	17.91	19.30	24.43	28.17	9.52
	+standard	17.86	12.71	22.99	19.40	24.36	29.37	31.08	9.52
	+local	52.50	40.49	57.02	95.52	36.24	33.23	37.23	80.95
	+standard+local	52.54	41.04	56.34	89.55	39.79	39.68	43.67	77.78
politics-government	question-only	28.96	20.33	41.86	12.50	13.26	17.88	22.96	20.00
	+standard	32.03	24.92	35.28	25.00	19.97	21.87	24.48	30.00
	+local	59.47	50.32	66.60	75.00	28.72	30.98	35.11	80.00
	+standard+local	59.26	48.81	63.18	62.50	30.20	35.87	41.66	60.00
education	question-only	38.95	28.78	45.86	20.00	29.06	34.44	39.73	33.33
	+standard	33.09	28.79	50.63	20.00	51.07	53.79	58.32	66.67
	+local	85.46	68.38	80.75	100.00	46.77	41.63	37.92	100.00
	+standard+local	79.88	62.84	76.39	80.00	57.71	62.53	66.00	100.00
transportation	question-only	19.59	14.72	25.05	0.00	6.96	14.41	14.87	0.00
	+standard	24.05	18.67	29.41	8.00	12.27	26.49	18.70	0.00
	+local	62.15	49.81	65.15	96.00	13.82	23.24	14.61	66.67
	+standard+local	64.52	53.78	63.45	84.00	21.45	20.79	16.70	100.00
food	question-only	6.50	8.01	23.13	0.00	28.65	33.24	42.23	50.00
	+standard	17.10	8.19	20.57	12.50	31.35	44.64	48.60	50.00
	+local	53.98	39.67	51.11	87.50	28.95	42.31	40.30	75.00
	+standard+local	53.84	39.42	51.00	87.50	27.87	43.94	42.66	50.00
linguistics	question-only	15.42	4.14	20.12	0.00	11.14	16.84	22.44	16.67
	+standard	17.04	7.01	19.84	16.67	16.99	24.82	23.37	25.00
	+local	45.30	25.83	48.70	83.33	31.12	37.85	43.94	83.33
	+standard+local	59.60	42.43	63.90	100.00	34.99	41.05	45.74	83.33
culture-customs	question-only	11.94	6.29	10.45	12.50	31.75	34.30	41.50	41.67
	+standard	20.49	7.71	13.62	12.50	23.61	31.18	32.85	33.33
	+local	51.93	38.79	50.74	75.00	50.32	52.24	55.31	100.00
	+standard+local	52.15	35.92	51.79	87.50	41.80	45.13	50.05	100.00
other	question-only	15.87	10.81	27.42	33.33	23.43	25.50	28.40	14.29
	+standard	12.37	8.90	25.60	33.33	31.99	40.09	38.61	28.57
	+local	40.02	26.36	36.28	66.67	37.59	44.14	47.21	85.71
	+standard+local	42.51	26.95	38.60	66.67	39.00	51.04	55.64	85.71
history	question-only	5.45	4.40	7.59	18.18	19.34	23.53	25.55	20.83
	+standard	10.05	5.96	5.70	9.09	20.95	26.83	26.74	16.67
	+local	52.97	38.51	48.87	90.91	35.34	40.31	45.13	66.67
	+standard+local	50.38	36.70	42.57	81.82	36.39	42.08	47.05	54.17
entertainment-art	question-only	21.19	13.04	24.00	18.75	15.06	20.27	22.79	16.22
	+standard	25.18	14.66	28.26	16.67	20.89	25.17	27.59	24.32
	+local	56.22	39.61	56.46	83.33	33.40	36.82	40.42	86.49
	+standard+local	57.81	42.81	57.18	79.17	35.99	40.87	44.07	83.78

Table 9: Summary-level QA results as measured by all metrics for **Qwen2.5-72B** across 11 topics.

Region	Context	yue-cmn				bar-deu			
		rouge-l	chrF++	BERTScore	LLMaJ	rouge-l	chrF++	BERTScore	LLMaJ
Local	question-only	18.81	13.03	26.04	9.90	19.01	22.62	26.86	16.67
	+standard	21.81	15.31	26.26	13.86	25.63	29.63	31.65	18.33
	+local	57.54	45.27	59.55	91.09	39.11	38.62	43.46	83.33
	+standard+local	59.02	47.02	58.57	83.17	42.00	44.02	49.19	76.67
	+local (translated)	60.65	47.76	59.34	89.11	/	/	/	/
Standard	question-only	17.76	12.96	23.64	13.79	20.77	29.09	32.16	27.27
	+standard	21.04	12.77	22.82	17.24	27.97	36.88	41.35	27.27
	+local	53.59	39.55	53.13	82.76	40.62	50.52	52.68	90.91
	+standard+local	53.41	39.09	52.23	82.76	35.17	46.90	50.19	81.82
	+local (translated)	56.35	40.99	53.18	79.31	/	/	/	/
Other	question-only	15.83	10.03	21.33	20.73	18.08	23.24	26.48	16.82
	+standard	17.94	11.13	22.51	25.61	20.85	27.04	27.29	20.56
	+local	49.18	34.67	54.05	87.20	32.16	34.41	37.22	79.44
	+standard+local	50.19	36.60	53.18	86.59	34.60	38.53	41.54	76.64
	+local (translated)	56.72	39.96	54.30	85.98	/	/	/	/

Table 10: Summary-level QA results as measured by all metrics for **Qwen2.5-72B**, divided by article region classification (Local, Standard, Other).