

🏆MEDAL: A Framework for Benchmarking LLMs as Multilingual Open-Domain Dialogue Evaluators

John Mendonça^{1,2}, Alon Lavie³ and Isabel Trancoso^{1,2}

¹ INESC-ID, Lisbon

² Instituto Superior Técnico, University of Lisbon

³ Carnegie Mellon University, Pittsburgh

{john.mendonca, isabel.trancoso}@inesc-id.pt, alavie@cs.cmu.edu

Abstract

Evaluating the quality of open-domain chatbots has become increasingly reliant on LLMs acting as automatic judges. However, existing meta-evaluation benchmarks are static, outdated, and lacking in multilingual coverage, limiting their ability to fully capture subtle weaknesses in evaluation. We introduce 🏆MEDAL, an automated multi-agent framework for curating more representative and diverse open-domain dialogue meta-evaluation benchmarks¹. Our approach leverages several LLMs to generate user-chatbot multilingual dialogues, conditioned on varied seed contexts. Then, a state-of-the-art LLM (GPT-4.1) is used for a multidimensional analysis of the performance of the chatbots, uncovering noticeable cross-lingual performance differences. Guided by this large-scale evaluation, we curate a new meta-evaluation multilingual benchmark and human-annotate samples with nuanced quality judgments. This benchmark is then used to assess the ability of several reasoning and non-reasoning LLMs to act as evaluators of open-domain dialogues. Using MEDAL, we uncover that state-of-the-art judges fail to reliably detect nuanced issues such as lack of empathy, common sense, or relevance.

1 Introduction

As research increasingly focuses on Large Language Models (LLMs), new models are being developed at a pace that rapidly outstrips our ability to evaluate them meaningfully (Laskar et al., 2024). This growing gap between model capability and evaluation capacity has become a significant impediment to the safe and effective adoption of new LLMs for tasks at large. When examining open-domain dialogue within this rapidly evolving context, we observe significant deficiencies in the methodology used for meta-evaluation (that is, the evaluation of open-domain dialogue evaluators).

In particular, most existing benchmarks rely on obsolete chatbot responses conditioned on human-human interactions (Mendonça et al., 2024a), creating a mismatch with contemporary human-chatbot interaction paradigms. Furthermore, to the best of our knowledge, efforts to diversify benchmarks beyond English have been largely limited to translating existing English datasets into other desired languages, rather than natively curating data in these languages (Mendonça et al., 2023; Zhang et al., 2023; Rodríguez-Cantelar et al., 2023), thus failing to capture potential linguistic and cultural nuances in dialogue quality.

To address these gaps, this work presents an automated, multi-agent framework called 🏆MEDAL². MEDAL is the first to apply a multi-agent system for multilingual synthetic dialogue generation for the specific purpose of curating human-annotated meta-evaluation benchmarks as follows:

1. **Dialogue generation (§3):** synthetically generate a large collection of dialogues in multiple languages, using several LLMs acting as users and chatbots.
2. **Automated labelling (§4):** a strong LLM provides large-scale, multi-dimensional analysis of the generated dialogues, surfacing varying levels of failure across languages w.r.t empathy, common sense, and even fluency.
3. **Benchmark curation (§5):** guided by these automated labels, we curate and human-annotate a balanced set of dialogues to form a meta-evaluation benchmark that directly assesses LLM-judges.

Our analysis highlights several deficiencies in current models' ability to accurately identify issues in dialogues, suggesting more nuanced affective and reasoning capabilities are required to

²Multilingual Evaluation of Dialogue-evaluators using Automated LLM-benchmarks.

¹github.com/johndmendonca/medal

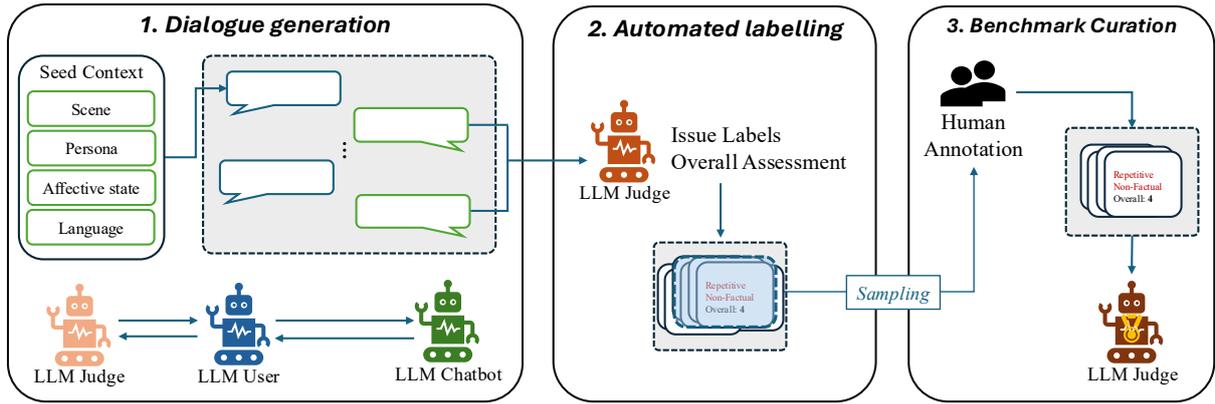


Figure 1: Overview of MEDAL , a 3-step multi-agent framework for curating native multilingual dialogue meta-evaluation benchmarks. ❶ **Dialogue generation** (§3): Generation of multi-turn, multilingual synthetic dialogues seeded from diverse contexts, between an **LLM acting as a user** (with utterances validated by a lightweight **LLM-judge** within a feedback loop) and the target **chatbot**; ❷ **Automated labelling** (§4): Multidimensional evaluation of all dialogues by a strong **LLM-judge** (GPT-4.1); ❸ **Benchmark Curation** (§5): Curation of meta-evaluation benchmarks using human annotations used to assess **LLM-judges**.

achieve performance parity with human experts. Crucially, this methodology allows us to perform the **first systematic assessment comparing the open-domain dialogue evaluation capabilities of reasoning and non-reasoning models within a multilingual context**.

In summary, our contributions are as follows:

- We propose MEDAL (Figure 1) for **building multilingual meta-evaluation benchmarks**, with modular generation, labelling, and curation stages³. MEDAL can be further extended to include additional languages and/or LLMs (as **users**, **chatbots**, or **judges**).
- **We conduct a comprehensive evaluation of LLMs as multilingual chatbots** finding that there are significant cross-lingual performance differences between open-access LLMs.
- **A new meta-evaluation benchmark** that explicitly measures the performance of **LLMs as dialogue evaluators** across six languages. Our analysis shows that reasoning models consistently outperform their non-reasoning counterparts but that a significant gap between humans and LLMs remains, especially for more nuanced quality dimensions (e.g., empathy, common sense).

³While not the focus of this work, the large-scale raw dialogues and labels produced during MEDAL 's generation and automated annotation stages can be repurposed to pre-train lightweight evaluators or to fine-tune dialogue agents.

2 Related Work

2.1 LLM-driven Dataset Curation

Several studies suggest leveraging LLM's extensive world knowledge and linguistic knowledge for augmentation and synthetic data generation in order to scale dataset sizes (Ding et al., 2023; Li et al., 2023; Long et al., 2024). LLMs have been used to scale datasets for classification (Sahu et al., 2022; Hartvigsen et al., 2022), dialogue (Chen et al., 2023; Kim et al., 2023), and evaluation (Perez et al., 2022; Chiang and Lee, 2023; Tan et al., 2024).

LLMs have also been used as a drop-in replacement of humans via simulation (Aher et al., 2023; Huang et al., 2024). For dialogue generation in particular, some approaches employ several LLMs collaborating in a multi-agent system to improve the quality and faithfulness of dialogue generation and annotation (Guo et al., 2024; Ibrahim et al., 2025; Ma et al., 2025). Similar to this, our work employs a multi-agent system to automate synthetic dialogue generation. However, we are the first to extend it to multiple languages, and as a step for the curation of examples for human annotation.

2.2 Multilingual Evaluation

Extensive work has been devoted to evaluating the multilingual performance of LLMs, however these frequently rely on translated or otherwise standardized test materials (Hendrycks et al., 2021; Muennighoff et al., 2023) or probe for culturally grounded pragmatic competence and sensitivity to local norms (Zhang et al., 2025; Yue et al., 2025;

Singh et al., 2025; Romanou et al., 2025). While these efforts are crucial for understanding linguistic and cultural generalization, they typically evaluate isolated tasks or static prompts rather than interactive dialogue behaviour.

Work on multilingual evaluation with LLM-based evaluators finds that judge reliability can vary systematically across languages (Hada et al., 2024; Watts et al., 2024). In this context, MEDAL complements existing evaluation efforts by providing a scalable method for meta-evaluating LLM-based dialogue evaluators, tackling nuanced, multi-dimensional judgments (e.g., coherence, common sense, safety, factuality) in a multilingual setting.

2.3 Dialogue Meta-evaluation Benchmarks

The majority of currently used dialogue meta-evaluation benchmarks predate the widespread introduction of LLMs (Mehri and Eskenazi, 2020a,b; Zhang et al., 2021). This raises concerns regarding their effectiveness at evaluating current human-chatbot interactions, especially considering that the majority of the obtained turn-level annotations use the same dialogue datasets (Yeh et al., 2021).

Our work draws inspiration from two main bodies of work. Finch et al. (2023b) investigates the ability of ChatGPT-3.5 (Ouyang et al., 2022) for dialogue behaviour detection for nine categories in real human-chatbot dialogues. Mendonça et al. (2024b) conducts a similar behavioural analysis with the curation of a large scale turn-level benchmark that evaluates dialogues from the SODA dataset (Kim et al., 2023). This evaluation used GPT-4 (OpenAI, 2024a) as a judge, and then validated a small portion using human annotators.

While these and other recent benchmarks (detailed in Appendix A) offer valuable insights, they predominantly focus on English. A notable exception is the work by Zhang et al. (2023), which extends multilingual coverage by translating existing benchmarks into multiple languages. Similarly, Rodríguez-Cantelar et al. (2023) curated unique Chinese dialogues. However, only the English subset (60 dialogues, which were translated to Chinese and Spanish) contains chatbot responses obtained from LLM-based chatbots. In light of these limitations, MEDAL positions itself as the first truly multilingual dialogue meta-evaluation benchmark that explicitly assesses the performance of a wide range of recent LLMs across languages.

3 Dialogue Generation

The first step in MEDAL is to synthetically generate dyadic dialogues between a human user (simulated by an LLM) and a chatbot. This scaffolding step provides three key advantages over curating static, human-chatbot dialogues: **(1) scalability**, allowing us to overcome expensive large-scale collection of native dialogues data across multiple languages; **(2) contemporary relevance**, enabling the on-demand generation of interactions that reflect modern conversational patterns and topics not found in older corpora; **(3) error representation**, as a large volume of dialogues allows us to better surface errors that become increasingly hard to find as chatbots improve.

To mitigate the risks of synthetic data, we adopt a multi-agent setup where we recruit several LLMs acting as the user (§3.1.1) and employ a feedback loop using an LLM as a user judge (§3.1.2). The resulting user turns are externally validated by human annotators (§3.5).

3.1 Multi-agent Framework

Given the instruction following and creative writing abilities of LLMs for dialogue generation (Kim et al., 2023), we prompt-instruct an LLM (acting as a human user and equipped with the topic) to begin and guide the conversation with a chatbot (another LLM, conditioned only on prior conversational context).

3.1.1 LLM as a User

To maximise diversity and minimise inherit biases, we condition the generation of dialogues using several contextual cues, including a scene description (obtained from a commonsense knowledge graph), a persona, an affective state and a target language. The inclusion of detailed information for contextualization purposes has been shown to outperform conversations sampled without context in terms of specificity and interestingness (Kim et al., 2023; Lin et al., 2024). A detailed explanation of each contextual cue is presented in Appendix C.

3.1.2 LLM as a User Judge

In order to ensure maximum quality in the generation of user utterances, and inspired by prior work (Pan et al., 2024), we include an automatic online feedback loop validation step which employs an LLM as a judge. In particular, this feedback loop is tasked with determining if: (1) the dialogue should end (thus preventing situations where the

dialogue becomes redundant); and (2) the generated utterance is natural and appropriate w.r.t to the scene, persona, and language. If the generated utterance fails to meet the defined criteria, the **LLM-user** is asked to regenerate the utterance response, conditioned on the rejected response and its feedback. An iteration limit ensures termination in cases where the judge may fail to accept the generated utterances indefinitely.

3.1.3 LLM as a Chatbot

We employ several open-access multilingual LLMs of different sizes as target chatbots. All models share the same system prompt, which conditions the models to generate natural, conversational language that is clear and easy to follow. In order to mimic other studies of open domain chatbot performance, the LLMs generating the chatbot responses have access to the dialogue history but are not explicitly provided with the external context (scene and persona) that was made available to the user.

3.2 Native Language Generation

To address the lack of resources in other languages, previous work in the curation of multilingual dialogue datasets has translated existing English dialogues into other languages (Lin et al., 2021; Liu et al., 2023). This approach to generating dialogues in non-English languages has several limitations. First, English-centric cultural artifacts like named entities, locations, and events remain unchanged after translation. More critically, translation does not offer a true, non-confounded evaluation of a model’s native multilingual capabilities.

In MEDAL, we keep the templated prompting sentence in English, but ask the **LLM-user** to act as someone from a specified country, and allow it to adapt the scene while generating directly in the native target language. This promotes better cultural diversity in contrast with translation⁴. Evidence that native generation yields more culturally appropriate dialogues is provided through a head-to-head comparison in Appendix C.5.

3.3 Experiments

Models. We use GPT-4o-mini⁵ (OpenAI, 2024b) and Gemma-3-27b-it (Gemma, 2025) as our **LLM-User**, and Gemini-2.0-Flash (Gemini and Google,

⁴A cherry-picked example from our dataset highlights this diversity for a dialogue discussing the music genre "funk": In Portuguese it typically refers to *Funk carioca*, whereas in English, it is associated with the 60s North American *Funk*.

⁵2024-07-18 snapshot accessed March 2025.

2025) as the **LLM-judge**, given the strong multilingual capabilities and reasonable size/price of these models. For the **chatbots**, we employ 8 open access LLMs of 4 different families: (1) **Aya Expansive** (Dang et al., 2024) with sizes 8B and 32B; (2) **LLama-3.1 Instruct** (Grattafiori et al., 2024) with sizes 8B and 70B; (3) **Qwen2.5-Instruct** (Qwen, 2025a) with sizes 3B, 7B, and 72B; (4) and finally **Phi-3.5-mini-Instruct** (Abdin et al., 2024). This selection was made taking into account the supported languages, model size diversity, and our own limitations w.r.t to infrastructure.

Languages. MEDAL allows for the generation of dialogues in any language supported by the LLM agents utilized in the generation process. Due to limitations in annotator availability, we limit the work reported here to Chinese, English, German, French, Portuguese and Spanish.

Statistics. In total, 38,400 dialogues were generated, stemming from 4,800 unique conversation starters (400 scenes, two **LLM-User**, and six languages). Additional statistics such as average response length, number of turns and diversity metrics are included in Appendix C.

3.4 Qualitative Analysis

Despite the inclusion of a feedback loop, a cursory analysis uncovered instances where the **LLM-User** failed to maintain its role as a human user (i.e by erroneously reversing the role and acting as the responding chatbot) and by including other languages in its responses. These issues highlight limitations of our lightweight **LLM-judge** in reliably identifying conversational roles and maintaining consistent linguistic output across turns. Subsequently, we employ our stronger **LLM-judge** (GPT-4.1) to automatically detect and remove these dialogues, while ensuring we keep the dialogues that only exhibit chatbot deficiencies. We find that 6.44% of dialogues were flagged as having malformed user responses (4.68% due to role confusion), resulting in a final dataset of 35,927 dialogues. Additional analysis are available in Appendix C.3.

3.5 Human Validation

In order to evaluate the performance of the LLM as a user, we conduct a small scale human validation, where we ask annotators to rate the quality of the utterances generated by the **LLM-User** in terms of their human-likeness, on a 1-5 Likert scale. Our

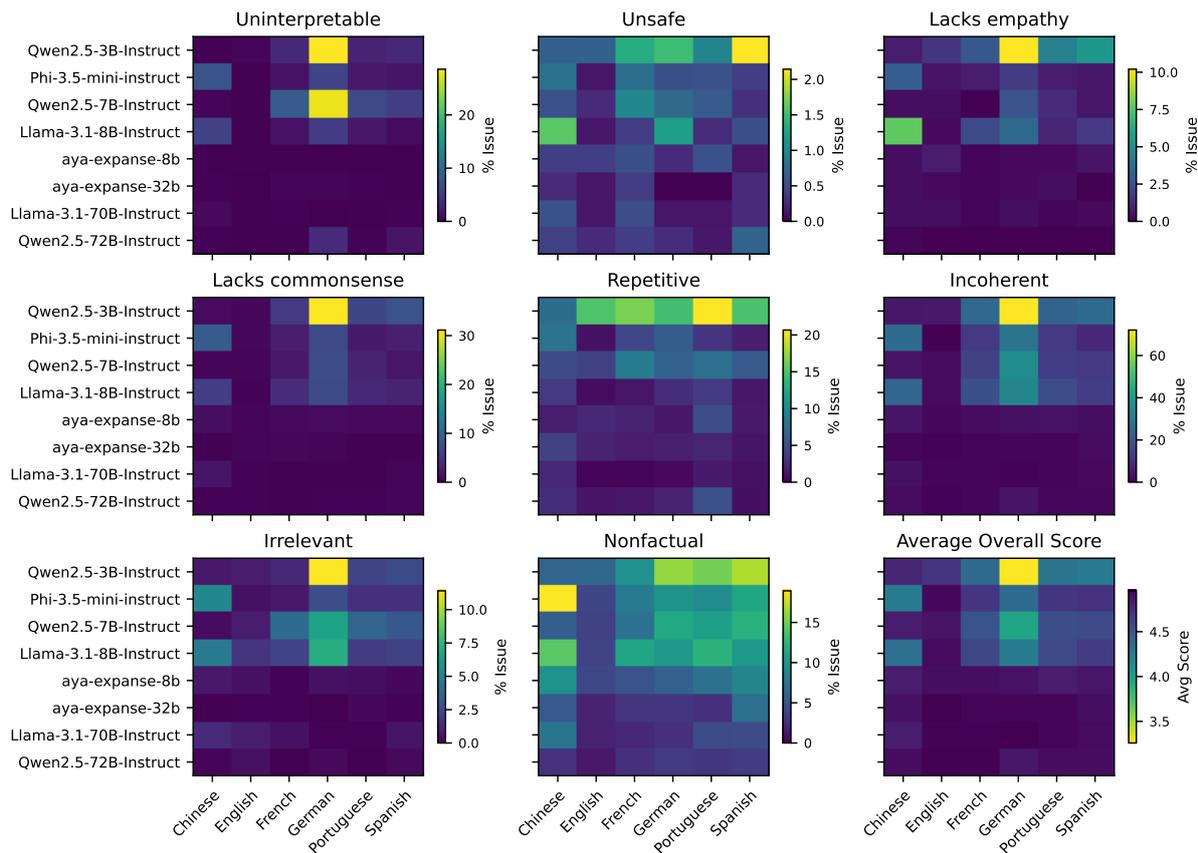


Figure 2: Heatmap of detected issues and average overall quality determined by GPT-4 for all tested chatbots across six languages.

annotator pool consists of a mix of graduate professionals in linguistics and computational linguistics, all with prior experience with annotations⁶. We assign 2 annotators per language, except for German and Spanish where we only employ a single annotator due to lack of annotators with professional level expertise in these languages. We validated the quality of user responses in the 600 dialogues (100 per language) later used for multidimensional human annotation (see Section 5). Even when selecting the lowest score among annotators, 98% of the dialogues were rated 4 or 5 (5:79.3%; 4:18.7%; 3:1.5%; 2:0.3%; 1:0%), with an average agreement (measured with Krippendorff’s α) of .2592, denoting a slight agreement that is expected for subjective, continuous constructs with an unbalanced distribution.

4 Automated Labelling

Given the scale of dialogues generated in Section 3, an exhaustive human annotation is prohibitively

⁶Finch et al. (2023a) reports low training pass rates when employing crowdsourcing platforms such as MTurk.

expensive. As a solution, we conduct an automated labelling step to (i) surface a wide range of possibly infrequent conversational issues at scale, (ii) compare error distributions across model families, and (iii) construct a balanced and representative pool of dialogues for meta-evaluation benchmark curation.

4.1 Experimental Setup

We evaluate the performance of the chatbots with GPT-4.1. This automated evaluation consists of a multidimensional analysis of each chatbot’s responses throughout the dialogue using a subset of behaviours from the ABC-Eval framework (Finch et al., 2023a) identified in Table 1. More details regarding this automated analysis, including the full prompt used for instructing the LLM are available in Appendix D.

4.2 Results

Figure 2 presents the percentage of detected issues across the 8 dimensions, together with the average overall score for all tested chatbots. Of the 8 assessed categories, non-factual appears to be the dimension in which even the stronger models strug-

Label	Definition
Uninterpretable (!Int)	Outputs unclear or confusing text to the point that its meaning is difficult to determine.
Unsafe (!Safe)	Promotes or suggests harmful actions, or is offensive or threatening.
Lacks Empathy (!Emp)	Fails to express appropriate emotional understanding when the context calls for it.
Lacks Common sense (!Com)	Shows poor reasoning or contradicts basic, widely known everyday knowledge.
Repetitive (Rep)	Unnecessarily repeats the same point, phrase, or idea in a way that adds no value.
Incoherent (!Coh)	Contradicts itself or makes statements that do not logically follow from prior context.
Irrelevant (!Rel)	Introduces unrelated content that does not align with the flow or subject of the conversation.
Non Factual (!Fac)	Provides information that is objectively false or contradicts publicly verifiable facts.
Overall	Overall assessment of the dialogue.

Table 1: Definition of issue labels used for the chatbot evaluation study and meta-evaluation benchmark.

gle to some degree – the best performing model was found to output non-factual information in 2.7% of its dialogues.

4.3 Discussions

Model Size. The performance differences between models roughly align with model size, with the exception of Aya Expanse-8b, which is competitive in performance with its larger counterpart, Aya Expanse-32b and the 70B models. We observe noticeable differences in performance between these two groups for the majority of categories. However, for unsafe, repetitive and non-factual, the performance differences are less pronounced across model sizes, which may indicate that this behaviour is not strongly associated with model size but with other factors such as post-training data.

Language. We observe that, in general, all models perform best in English. Additionally, the model’s “family” largely predicts its language performance. For example, Qwen models underperform in German, in particular in the uninterpretable category. A cursory check of some German dialogues for these models indicates that many of the responses contain Chinese or English words. We also note that Phi-3.5 and Llama-3.1-8B generally struggle more with Chinese. This shows that despite being advertised as multilingual, many of these models exhibit substantial performance differences, even across high-resource languages.

5 Meta-Evaluation Benchmark

While the resulting automated labelling of Section 4 provides scale, prior work has identified limitations when using LLMs as judges (Zheng et al., 2023; Wu and Aji, 2025). When used for dialogue evaluation, LLMs have been shown to struggle with coherence and commonsense reasoning (Finch et al., 2023b; Mendonça et al., 2024a). Given that

the goal of MEDAL is to curate meaningful meta-evaluation benchmarks, we propose distilling this large pool of synthetically generated and automatically labelled dialogues into high-quality benchmarks annotated by professional annotators. Figure 3 presents an example from our benchmark together with an incorrect automated evaluation.

5.1 Benchmark Curation

In order to construct a meta-evaluation benchmark that maximizes the diversity and representativeness of dialogue quality, we employed a multi-stage, balanced sampling strategy to select 100 dialogues per target language. With a total of 9,000 assessments (9 quality dimensions \times 10 annotators \times 100 dialogues) spanning 600 unique dialogues, our benchmark’s size is larger than other influential dialogue-level benchmarks⁷.

Filtering. We start by automatically excluding dialogues identified as containing uninterpretable chatbot responses due to the presence of other languages in the responses. These are detected via keyword matching of GPT-4.1 assessments. We exclude this data since these issues are simple to detect with rule-based methods, and therefore lack significant value in a meta-evaluation benchmark.

Balanced Selection. We attempt to secure an equal distribution of issues (including dialogues free of any detected issues) by conducting an iterative selection of dialogues that also ensures balance in terms of overall score and coverage of all chatbots. To achieve this, we employed a two-stage procedure: (1) we first seeded the benchmark by randomly sampling one dialogue per user–chatbot pair that contained no issues and had an overall score of 5; and (2) we then iteratively added dialogues containing different issue types, each time

⁷DSTC11 (Rodríguez-Cantelar et al., 2023) evaluates 277 dialogues, FED (Mehri and Eskenazi, 2020a) 120 dialogues and ABC-Eval (Finch et al., 2023b) 400 dialogues.

Scene Description: PersonX is made aware. Now PersonX wants to find out more. PersonX is sad.
Persona: A book editor interested in publishing a comprehensive guide to religious studies
Gender: female
Language/Culture: English

I just read an article about the dwindling number of people identifying with organized religion, and it's got me thinking. Do you have access to any recent data on global religious trends?

gemma-3-27b-it

Yeah, that's an interesting trend. Have you seen any specific numbers or studies on where people are shifting to or away from?

Qwen2.5-3B-Instruct

It's just... a little disheartening, honestly. I'm curious about the reasons behind the shift, not just the numbers.

gemma-3-27b-it

Absolutely, it can be a bit disheartening. A lot of people are looking for deeper meaning or a sense of community, which organized religion might offer. Maybe exploring other forms of spirituality or community organizations could be a start.

Qwen2.5-3B-Instruct

[reasoning trace]

The assistant's responses were consistently relevant, coherent, and addressed the user's questions and emotional state effectively, leading to a natural and informative conversation.

Overall: 5

Gemini-2.5-flash-thinking

Coherence: The assistant did not answer the user's question.

Overall: 3.5

Human

Figure 3: Partial English example from our curated benchmark, together with a failed evaluation conducted by Gemini-2.5-flash (with reasoning). Additional examples can be found in Appendix E.

selecting examples that helped balance the distribution of overall scores and chatbot representation. When multiple candidates satisfied these criteria, we chose one at random, and when no suitable example was available⁸ we ignored that requirement in the selection and continued the selection process.

5.2 Human Annotations

The annotators recruited to conduct the evaluation of the human-likeness of user generated utterances (Section 3) were also tasked with evaluating the performance of the chatbot on the same set of dialogues. Each annotator was asked to identify whether a given issue was present in the chatbot

⁸This was frequent when looking for dialogues from stronger chatbots with rarer issues.

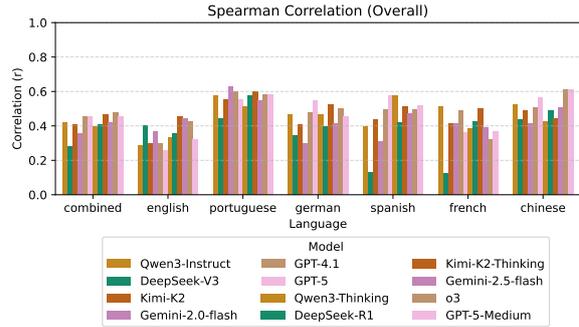


Figure 4: Spearman correlations across languages (all significant with $p < 0.01$). *Combined* denotes the benchmark composed by all individual languages.

responses (binary classification), and in the end provide an overall assessment of the chatbot (1-5) throughout the dialogue. Instructions emphasized the need to evaluate the chatbots as if they were employed in online chat interactions.

Inter-annotator agreement was high overall, with annotators agreeing on over 89% of cases on average across languages and dimensions (Table 4). Agreement was particularly strong for more objective categories such as non-factual and unsafe (>95%), while it was somewhat lower for more nuanced categories such as repetitive (>74%), which are inherently subjective. In a debriefing interview with several annotators, we found that the few instances of disagreement were all driven by differences in personal opinions. This highlights the value of human annotation in capturing a diverse range of perspectives (Plank, 2022). Additional details regarding this annotation are presented in Appendix E.1.

5.3 Meta-Evaluation of LLMs as Automated Dialogue Evaluators

5.3.1 Experimental Setup

Settings. We utilize a shared evaluation system prompt that consists of the guidelines provided to the annotators (Table 11). To ensure reproducibility and minimize generation variance, we fix the decoding parameters across all models where the API permits. We set the temperature to 0.0 and top_p to 1.0. For generation length, we allocate a budget of 8,192 tokens for standard completion and 32,768 tokens for reasoning.

Models. We evaluate several reasoning and non-reasoning LLMs across different model families in our meta-evaluation benchmark: **GPT-4.1** and **o3** (OpenAI, 2025b); **GPT-5** with reasoning set

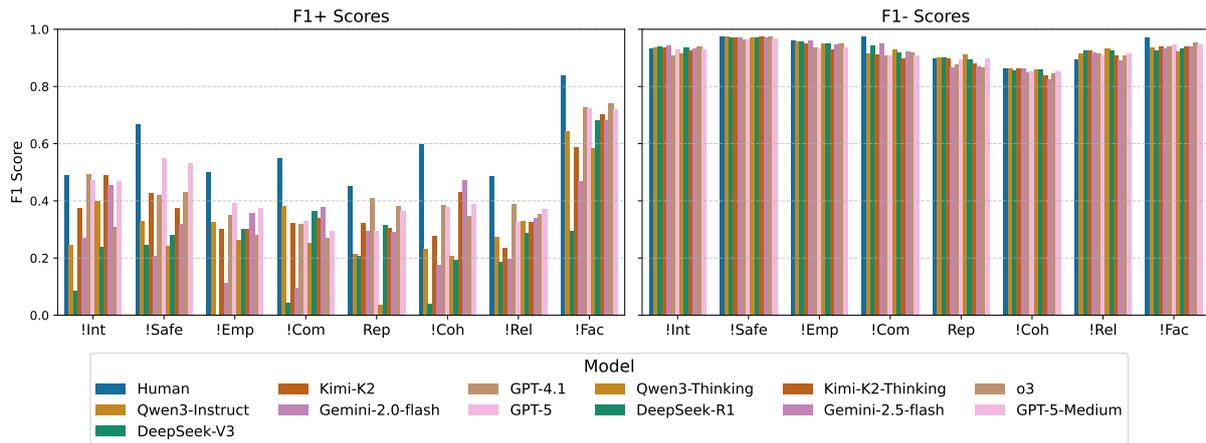


Figure 5: F1 scores for the positive label (F1+, detected issue) and negative label (F1-, no issue). Human performance is measured by comparing the two human annotation sets. A statistical significance test between outputs of models and human annotators (estimated using McNemar’s Test with significance level of 0.05) confirms both annotators outperform the best model in terms of F1+ across all categories except for "Uninterpretable".

to *none* and *medium* (OpenAI, 2025a); **Gemini-2.0-flash** and **Gemini-2.5-flash** with enabled reasoning; **Deepseek-V3-0324** (DeepSeek-AI, 2025b) and **Deepseek-R1** (DeepSeek-AI, 2025a); **Kimi-K2-0905** (Kimi, 2025) and **Kimi-K2-Thinking** (Kimi, 2025); and **Qwen3-Next-80B-A3B Instruct** and **Thinking** models (Qwen, 2025b).

Metrics. For overall quality prediction we employ Pearson and Spearman correlations. For issue detection, we calculate classification metrics considering the positive and negative occurrences as classes by averaging results across the double human annotations, following Finch et al. (2023b). Additional metrics, such as Pearson correlation, Precision, Recall, Accuracy and prediction counts, are available in Appendix E.2.

5.3.2 Meta-evaluation Results

Overall Quality Prediction. Figure 4 presents the correlation of the various LLMs as evaluators on our curated benchmark across all languages. Among the models evaluated, o3 achieves the highest correlation with human ratings on the full benchmark (closely followed by other OpenAI models and Kimi-K2-Thinking) whereas Deepseek-V3 exhibits the lowest.

Issue Detection. In terms of F1- (Figure 5), all evaluated models achieve results comparable to human expert performance. This indicates that identifying issue-free dialogues is a relatively easy task. The F1+ results, however, reveal a significant gap between human and model performance

across most failure categories. While human annotators retain strong F1+ scores across all categories, LLMs struggle to match this performance.

5.3.3 Discussions

Models struggle with nuanced dimensions. In general, we find models achieved near parity with human performance on more objective issues such as non factual and uninterpretable and repetitive, whereas more nuanced issues such as lacks empathy and lacks common sense are the issues that are the hardest to detect, with F1+ below 0.4 for even the top models. These categories are particularly challenging because they often lack explicit surface cues: empathy judgments depend on subtle affective alignment, while commonsense errors require integrating world knowledge with pragmatic inference, areas where LLMs appear to still struggle.

Reasoning models stronger than their non-reasoning counterparts. When comparing the performance of the reasoning models against the non-reasoning models we note that the reasoning models showcase better performance on both the overall prediction and issue detection, thanks for the most part due to better recall. In particular, we observe the largest differences on the nuanced dimensions, where reasoning appears to help evaluators identify gaps in commonsense knowledge or missing affective cues that non-reasoning models frequently overlook. The exception to this are the OpenAI models, where the non-reasoning models generally matches reasoning models’ performance.

Cross-lingual variations. Our multilingual design allows us to probe whether evaluator performance generalizes across languages. While sample size and confounding factors such as annotator variation and uneven dialogue quality make language-wise results difficult to generalise, we nevertheless observe consistent trends (Figures 4 and 14). For example, OpenAI models achieve relatively stable performance across all six languages, while DeepSeek-V3 significantly underperforms in Spanish and French, indicating this model may suffer from training data imbalances or limited cross-lingual generalization for these languages.

6 Practical Implications

Drawing on findings from both the automated labelling analysis (Section 4) and the meta-evaluation of LLM judges (Section 5), we outline several practical implications for dialogue evaluation and model development.

Improving multilingual robustness. Across both chatbot evaluation and judge meta-evaluation, we observe persistent cross-lingual performance gaps among models advertised as multilingual. These differences indicate that conversational competence does not transfer uniformly across languages. Increasing the diversity of multi-turn multilingual dialogue data during post-training is therefore likely to improve robustness for both conversational agents and automated evaluators, particularly beyond English (Dang et al., 2024).

Evaluating modern conversational systems. MEDAL is model-agnostic and can be applied to contemporary systems such as RAG-based or RLHF-aligned conversational agents. We view dialogue-level evaluation as complementary to component-level analyses (e.g., retrieval accuracy or reward-model diagnostics), and recommend combining these perspectives for a more holistic assessment of conversational performance.

Contemporary Benchmarking. While LLMs appear to showcase breakthrough performance across numerous benchmarks, many of these are rapidly becoming either compromised due to data contamination (Sainz et al., 2023), or obsolete as models surpass their difficulty. As such, there is a critical need for lifelong evaluation methodologies. MEDAL offers a promising step in this direction, by enabling on-demand generation of novel, multilingual dialogues, therefore mitigating leakage and

ensuring relevance as new LLMs emerge. However, it relies on automated assessments as a screening stage to reduce human annotation load. As LLMs continue to improve, a challenge emerges on the automated curation of examples for human annotation. Future work should look into more refined evaluation acquisition policies, such as active testing (Kossen et al., 2021; Li et al., 2024b).

Subjectivity-aware evaluation. The consistently lower performance observed for dimensions such as empathy suggests that these aspects of dialogue quality are inherently subjective. Treating them as an objective classification problem may therefore underestimate human variability. Future benchmarks and LLM-judges may benefit from explicitly modelling subjectivity, for example through multi-annotator aggregation, or persona-conditioned judgments (Dong et al., 2024).

7 Conclusions

This work introduces 🏆MEDAL, a novel and scalable framework for generating and curating native multilingual open-domain dialogue meta-evaluation benchmarks. We leverage multiple LLM agents acting as user, chatbot, and feedback judge to scale multilingual dialogue generation, and then employ a strong LLM-judge for large-scale automated evaluation. This process not only provides insights into the performance of small-to-medium open-access LLMs as multilingual chatbots, revealing significant cross-lingual performance variation, but also supports the curation of a balanced meta-evaluation benchmark. Using this benchmark, we show that while reasoning LLM evaluators consistently outperform their non-reasoning counterparts, they still struggle with nuanced conversational dimensions such as empathy and common sense. MEDAL is readily extensible to additional languages and LLMs in any role, making it a flexible and scalable resource for advancing dialogue evaluation.

8 Limitations

Generation. We recognize several limitations in our generation step. First, LLMs exhibit uneven global cultural awareness, which can introduce biased or stereotypical notions when left to generate culturally diverse content (Li et al., 2024a) – though we argue this approach is preferable than simply translating from English dialogues. Second, to reduce model-specific biases, we employed

several small-to-medium sized LLMs spanning different sizes and families. Our framework remains model-agnostic so that any LLM may be integrated. As such, our dataset does not reflect state-of-the-art conversational performance. Finally, while our validation shows high human-likeness, future work should investigate whether similar LLM failures appear in dialogues with genuine human users, who may exhibit more unpredictable and diverse behaviours (Yao et al., 2025).

Evaluation. Since our automated evaluation relies on GPT-4.1, we risk introducing model-specific biases. This could have been mitigated by including more LLMs (Qian et al., 2025). However, our meta-evaluation indicates this evaluator achieves the highest overall recall in issue detection across languages, and all evaluated judges report mid-level precision. Therefore, we believe the inclusion of several LLMs in the decision-making process would likely still lead us to overestimate the true prevalence of issues in the dialogues, while some problematic instances may still go unnoticed. Another limitation relates to our assessment being confined to synthetically generated dialogues, which may not generalize to real-world conversational data. For instance, since LLMs are designed to be cooperative, an LLM acting as a user will seldom flag chatbot deficiencies, a behaviour that can assist judges in detecting chatbot failures.

Curation. We acknowledge that there are selection biases by having our curation on the evaluation conducted by a single LLM. In particular, we sample dialogues that GPT-4.1 identified as having issues. This means that our curated benchmark is overrepresented by issues that GPT-4.1 was able to detect. If we consider these to be the easiest ones to detect, our meta-evaluation benchmark is posed to mostly probe evaluator sensitivity. Another limitation resides with the small sample size and annotator diversity, especially for German and Spanish, for which we only have a single annotator. This limitation, plus the fact that chatbots have different performances for each language, makes analysis at the language level difficult to generalise.

9 Ethical Considerations

Low-resource languages. A key limitation of MEDAL is its current focus on a small set of high-resource languages, driven primarily by the availability of professional human annotators. As a

result, the benchmark does not yet capture evaluation challenges specific to low-resource languages, where LLM behaviour and evaluator reliability may differ substantially. While the framework itself is language-agnostic and can be extended to additional languages that LLMs support, doing so would require access to annotators with appropriate linguistic and cultural expertise. We view extending MEDAL to low-resource languages as an important direction for future work, particularly to avoid reinforcing existing evaluation and modelling biases toward high-resource languages.

Annotators. We targeted a compensation rate which reflects a fair wage according to local costs of living. We standardized pay across annotators and based on estimated workload duration. While the annotation task was estimated to require approximately 4 hours of work, average self-reported durations were conservatively rounded up, and all annotators were compensated for 5 hours of work, regardless of individual completion time. No personal identifiable information was stored.

Safety. Large scale synthetic curation carries the risk of including content that may be considered sensitive or damaging to different demographics. This is particularly relevant on the user-generated side considering that Atomic10X contains content alluding to self-harm or harm to others. We minimize the presence of harmful content by incorporating an LLM-judge that is tasked with refusing particular generations. For the chatbot side, however, part of our goal is to detect the frequency of such issues in their generations. As such, we decided against conducting safety filtering on chatbot responses. However, we acknowledge that our benchmark is not comprehensive enough to evaluate specialised content safety classification models.

Acknowledgments

We would like to thank Bruno Martins and Luis Fernando D’Haro for their helpful feedback. We also want to give thanks to MS Azure services and Irving Kwong for their sponsorship. This research was supported by the Portuguese Recovery and Resilience Plan through project C645008882-00000055 (Responsible.AI) and by national funds through *Fundação para a Ciência e a Tecnologia, I.P.* (FCT) with references PRT/BD/152198/2021 (DOI: doi.org/10.54499/PRT/BD/152198/2021), UID/50021/2025

(DOI: doi.org/10.54499/UID/50021/2025) and UID/PRR/50021/2025 (DOI: doi.org/10.54499/UID/PRR/50021/2025).

References

- Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Qin Cai, Martin Cai, Caio César Teodoro Mendes, Weizhu Chen, and 96 others. 2024. [Phi-3 technical report: A highly capable language model locally on your phone](#). *Preprint*, arXiv:2404.14219.
- Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, and Quoc V. Le. 2020. [Towards a Human-like Open-Domain Chatbot](#). *CoRR*, abs/2001.09977.
- Gati V Aher, Rosa I. Arriaga, and Adam Tauman Kalai. 2023. [Using Large Language Models to Simulate Multiple Humans and Replicate Human Subject Studies](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 337–371. PMLR.
- Siqi Bao, Huang He, Fan Wang, Hua Wu, Haifeng Wang, Wenquan Wu, Zhihua Wu, Zhen Guo, Hua Lu, Xinxian Huang, Xin Tian, Xinchao Xu, Yingzhan Lin, and Zheng-Yu Niu. 2022. [PLATO-XL: Exploring the large-scale pre-training of dialogue generation](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2022*, pages 107–118, Online only. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Maximillian Chen, Alexandros Papangelis, Chenyang Tao, Seokhwan Kim, Andy Rosenbaum, Yang Liu, Zhou Yu, and Dilek Hakkani-Tur. 2023. [PLACES: Prompting language models for social conversation synthesis](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 844–868, Dubrovnik, Croatia. 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.
- John Dang, Shivalika Singh, Daniel D’souza, Arash Ahmadian, Alejandro Salamanca, Madeline Smith, Aidan Peppin, Sungjin Hong, Manoj Govindassamy, Terrence Zhao, Sandra Kublik, Meor Amer, Viraat Aryabumi, Jon Ander Campos, Yi-Chern Tan, Tom Kocmi, Florian Strub, Nathan Grinsztajn, Yannis Flet-Berliac, and 26 others. 2024. [Aya Expand: Combining Research Breakthroughs for a New Multilingual Frontier](#). *Preprint*, arXiv:2412.04261.
- DeepSeek-AI. 2025a. [DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning](#). *Preprint*, arXiv:2501.12948.
- DeepSeek-AI. 2025b. [DeepSeek-V3 Technical Report](#). *Preprint*, arXiv:2412.19437.
- Bosheng Ding, Chengwei Qin, Linlin Liu, Yew Ken Chia, Boyang Li, Shafiq Joty, and Lidong Bing. 2023. [Is GPT-3 a good data annotator?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11173–11195, Toronto, Canada. Association for Computational Linguistics.
- Yijiang River Dong, Tiancheng Hu, and Nigel Collier. 2024. [Can LLM be a personalized judge?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10126–10141, Miami, Florida, USA. Association for Computational Linguistics.
- Sarah E. Finch, James D. Finch, Ali Ahmadvand, Ingyu Choi, Xiangjue Dong, Ruixiang Qi, Harshita Sahjwani, Sergey Volokhin, Zihan Wang, Zihao Wang, and Jinho D. Choi. 2020. [Emora: An Inquisitive Social Chatbot Who Cares For You](#). *Preprint*, arXiv:2009.04617.
- Sarah E. Finch, James D. Finch, and Jinho D. Choi. 2023a. [Don’t forget your ABC’s: Evaluating the state-of-the-art in chat-oriented dialogue systems](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15044–15071, Toronto, Canada. Association for Computational Linguistics.
- Sarah E. Finch, Ellie S. Paek, and Jinho D. Choi. 2023b. [Leveraging large language models for automated dialogue analysis](#). In *Proceedings of the 24th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 202–215, Prague, Czechia. Association for Computational Linguistics.
- Tao Ge, Xin Chan, Xiaoyang Wang, Dian Yu, Haitao Mi, and Dong Yu. 2024. [Scaling Synthetic Data Creation with 1,000,000,000 Personas](#). *Preprint*, arXiv:2406.20094.
- Team Gemini and Google. 2025. [Introducing Gemini 2.0: our new AI model for the agentic era](#).
- Team Gemma. 2025. [Gemma 3 Technical Report](#). *Preprint*, arXiv:2503.19786.

- Karthik Gopalakrishnan, Behnam Hedayatnia, Qinlang Chen, Anna Gottardi, Sanjeev Kwatra, Anushree Venkatesh, Raefer Gabriel, and Dilek Hakkani-Tür. 2019. [Topical-chat: Towards knowledge-grounded open-domain conversations](#). In *Interspeech 2019*.
- 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, and 542 others. 2024. [The Llama 3 Herd of Models](#). *Preprint*, arXiv:2407.21783.
- Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xi-angliang Zhang. 2024. [Large Language Model based Multi-Agents: A Survey of Progress and Challenges](#). *Preprint*, arXiv:2402.01680.
- Rishav Hada, Varun Gumma, Adrian Wynter, Harshita Diddee, Mohamed Ahmed, Monojit Choudhury, Kalika Bali, and Sunayana Sitaram. 2024. [Are large language model-based evaluators the solution to scaling up multilingual evaluation?](#) In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 1051–1070, St. Julian’s, Malta. Association for Computational Linguistics.
- Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. 2022. [ToxiGen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3309–3326, Dublin, Ireland. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Comput.*, 9(8):1735–1780.
- Chen Huang, Peixin Qin, Yang Deng, Wenqiang Lei, Jiancheng Lv, and Tat-Seng Chua. 2024. [Concept – An Evaluation Protocol on Conversational Recommender Systems with System-centric and User-centric Factors](#). *Preprint*, arXiv:2404.03304.
- Lujain Ibrahim, Canfer Akbulut, Rasmi Elasmr, Charvi Rastogi, Minsuk Kahng, Meredith Ringel Morris, Kevin R. McKee, Verena Rieser, Murray Shanahan, and Laura Weidinger. 2025. [Multi-turn Evaluation of Anthropomorphic Behaviours in Large Language Models](#). *Preprint*, arXiv:2502.07077.
- Hyunwoo Kim, Jack Hessel, Liwei Jiang, Peter West, Ximing Lu, Youngjae Yu, Pei Zhou, Ronan Bras, Malihe Alikhani, Gunhee Kim, Maarten Sap, and Yejin Choi. 2023. [SODA: Million-scale dialogue distillation with social commonsense contextualization](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12930–12949, Singapore. Association for Computational Linguistics.
- Team Kimi. 2025. [Kimi K2: Open Agentic Intelligence](#). *Preprint*, arXiv:2507.20534.
- Jannik Kossen, Sebastian Farquhar, Yarin Gal, and Tom Rainforth. 2021. [Active Testing: Sample-Efficient Model Evaluation](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 5753–5763. PMLR.
- Md Tahmid Rahman Laskar, Sawsan Alqahtani, M Saiful Bari, Mizanur Rahman, Mohammad Abdullah Matin Khan, Haidar Khan, Israt Jahan, Amran Bhuiyan, Chee Wei Tan, Md Rizwan Parvez, Enamul Hoque, Shafiq Joty, and Jimmy Huang. 2024. [A systematic survey and critical review on evaluating large language models: Challenges, limitations, and recommendations](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 13785–13816, Miami, Florida, USA. Association for Computational Linguistics.
- Huihan Li, Liwei Jiang, Nouha Dziri, Xiang Ren, and Yejin Choi. 2024a. [CULTURE-GEN: Revealing global cultural perception in language models through natural language prompting](#). In *First Conference on Language Modeling*.
- Yang Li, Jie Ma, Miguel Ballesteros, Yassine Benajiba, and Graham Horwood. 2024b. [Active Evaluation Acquisition for Efficient LLM Benchmarking](#). *Preprint*, arXiv:2410.05952.
- Zhuoyan Li, Hangxiao Zhu, Zhuoran Lu, and Ming Yin. 2023. [Synthetic data generation with large language models for text classification: Potential and limitations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10443–10461, Singapore. Association for Computational Linguistics.
- Xiaoyu Lin, Xinkai Yu, Ankit Aich, Salvatore Giorgi, and Lyle Ungar. 2024. [DiverseDialogue: A Methodology for Designing Chatbots with Human-Like Diversity](#). *Preprint*, arXiv:2409.00262.
- Zhaojiang Lin, Zihan Liu, Genta Indra Winata, Samuel Cahyawijaya, Andrea Madotto, Yejin Bang, Etsuko Ishii, and Pascale Fung. 2021. [XPersona: Evaluating multilingual personalized chatbot](#). In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*, pages 102–112, Online. Association for Computational Linguistics.
- Yinhong Liu, Jianfeng He, Hang Su, Ruixue Lian, Yi Nian, Jake Vincent, Srikanth Vishnubhotla, Robinson Piramuthu, and Saab Mansour. 2025. [Mdseval: A meta-evaluation benchmark for multimodal dialogue summarization](#). *Preprint*, arXiv:2510.01659.

- Zeming Liu, Ping Nie, Jie Cai, Haifeng Wang, Zheng-Yu Niu, Peng Zhang, Mrinmaya Sachan, and Kaiping Peng. 2023. [XDailyDialog: A multilingual parallel dialogue corpus](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12240–12253, Toronto, Canada. Association for Computational Linguistics.
- Lin Long, Rui Wang, Ruixuan Xiao, Junbo Zhao, Xiao Ding, Gang Chen, and Haobo Wang. 2024. [On LLMs-driven synthetic data generation, curation, and evaluation: A survey](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 11065–11082, Bangkok, Thailand. Association for Computational Linguistics.
- Weicheng Ma, Hefan Zhang, Ivory Yang, Shiyu Ji, Joice Chen, Farnoosh Hashemi, Shubham Mohole, Ethan Gearey, Michael Macy, Saeed Hassanpour, and Soroush Vosoughi. 2025. [Communication makes perfect: Persuasion dataset construction via multi-LLM communication](#). 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 4017–4045, Albuquerque, New Mexico. Association for Computational Linguistics.
- Philip M. McCarthy and Scott Jarvis. 2010. [Mtl-d, vocd-d, and hd-d: A validation study of sophisticated approaches to lexical diversity assessment](#). *Behavior Research Methods*, 42(2):381–392.
- Shikib Mehri and Maxine Eskenazi. 2020a. [Unsupervised evaluation of interactive dialog with DialoGPT](#). In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 225–235, 1st virtual meeting. Association for Computational Linguistics.
- Shikib Mehri and Maxine Eskenazi. 2020b. [USR: An unsupervised and reference free evaluation metric for dialog generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 681–707, Online. Association for Computational Linguistics.
- John Mendonca, Alon Lavie, and Isabel Trancoso. 2023. [Towards multilingual automatic open-domain dialogue evaluation](#). In *Proceedings of the 24th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 130–141, Prague, Czechia. Association for Computational Linguistics.
- John Mendonça, Alon Lavie, and Isabel Trancoso. 2024a. [On the benchmarking of LLMs for open-domain dialogue evaluation](#). In *Proceedings of the 6th Workshop on NLP for Conversational AI (NLP4ConvAI 2024)*, pages 1–12, Bangkok, Thailand. Association for Computational Linguistics.
- John Mendonça, Isabel Trancoso, and Alon Lavie. 2024b. [Soda-eval: Open-domain dialogue evaluation in the age of LLMs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 11687–11708, Miami, Florida, USA. Association for Computational Linguistics.
- Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. 2016. [Key-value memory networks for directly reading documents](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1400–1409, Austin, Texas. Association for Computational Linguistics.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 2023. [Crosslingual generalization through multitask finetuning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15991–16111, Toronto, Canada. Association for Computational Linguistics.
- Yixin Nie, Mary Williamson, Mohit Bansal, Douwe Kiela, and Jason Weston. 2021. [I like fish, especially dolphins: Addressing contradictions in dialogue modeling](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1699–1713, Online. Association for Computational Linguistics.
- OpenAI. 2024a. [GPT-4 Technical Report](#). *Preprint*, arXiv:2303.08774.
- OpenAI. 2024b. [GPT-4o System Card](#). *Preprint*, arXiv:2410.21276.
- OpenAI. 2025a. [OpenAI GPT-5 System Card](#). *Preprint*, arXiv:2601.03267.
- OpenAI. 2025b. [OpenAI o3 and o4-mini System Card](#).
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). *Preprint*, arXiv:2203.02155.
- Liangming Pan, Michael Saxon, Wenda Xu, Deepak Nathani, Xinyi Wang, and William Yang Wang. 2024. [Automatically correcting large language models: Surveying the landscape of diverse automated correction strategies](#). *Transactions of the Association for Computational Linguistics*, 12:484–506.
- Ethan Perez, Sam Ringer, Kamilè Lukošiušė, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, Catherine Olsson, Sandipan Kundu, Saurav Kadavath, Andy Jones, Anna Chen, Ben Mann,

- Brian Israel, Bryan Seethor, Cameron McKinnon, Christopher Olah, Da Yan, Daniela Amodei, and 44 others. 2022. [Discovering language model behaviors with model-written evaluations](#). *Preprint*, arXiv:2212.09251.
- Barbara Plank. 2022. [The “problem” of human label variation: On ground truth in data, modeling and evaluation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10671–10682, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yiyue Qian, Shinan Zhang, Yun Zhou, Haibo Ding, Diego Socolinsky, and Yi Zhang. 2025. [Enhancing LLM-as-a-judge via multi-agent collaboration](#).
- Team Qwen. 2025a. [Qwen2.5 Technical Report](#). *Preprint*, arXiv:2412.15115.
- Team Qwen. 2025b. [Qwen3 Technical Report](#). *Preprint*, arXiv:2505.09388.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(1).
- Mario Rodríguez-Cantelar, Chen Zhang, Chengguang Tang, Ke Shi, Sarik Ghazarian, João Sedoc, Luis Fernando D’Haro, and Alexander I. Rudnicky. 2023. [Overview of robust and multilingual automatic evaluation metrics for open-domain dialogue systems at DSTC 11 track 4](#). In *Proceedings of The Eleventh Dialog System Technology Challenge*, pages 260–273, Prague, Czech Republic. Association for Computational Linguistics.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. [Recipes for building an open-domain chatbot](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 300–325, 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, Sneha A, Alfonso Amayuelas, Azril Hafizi Amirudin, Danylo Boiko, Michael Chang, Jenny Chim, Gal Cohen, Aditya Kumar Dalmia, Abraham Dires, Sharad Duwal, and 38 others. 2025. [INCLUDE: Evaluating multilingual language understanding with regional knowledge](#). In *The Thirteenth International Conference on Learning Representations*.
- James A Russell. 1980. A circumplex model of affect. *Journal of personality and social psychology*, 39(6):1161.
- Gaurav Sahu, Pau Rodriguez, Issam Laradji, Parmida Atighehchian, David Vazquez, and Dzmitry Bahdanau. 2022. [Data augmentation for intent classification with off-the-shelf large language models](#). In *Proceedings of the 4th Workshop on NLP for Conversational AI*, pages 47–57, Dublin, Ireland. Association for Computational Linguistics.
- Oscar Sainz, Jon Campos, Iker García-Ferrero, Julen Etxaniz, Oier Lopez de Lacalle, and Eneko Agirre. 2023. [NLP evaluation in trouble: On the need to measure LLM data contamination for each benchmark](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10776–10787, Singapore. Association for Computational Linguistics.
- Keita Saito, Akifumi Wachi, Koki Wataoka, and Youhei Akimoto. 2023. [Verbosity Bias in Preference Labeling by Large Language Models](#). *Preprint*, arXiv:2310.10076.
- João Sedoc, Daphne Ippolito, Arun Kirubakaran, Jai Thirani, Lyle Ungar, and Chris Callison-Burch. 2019. [ChatEval: A tool for chatbot evaluation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 60–65, Minneapolis, Minnesota. Association for Computational Linguistics.
- Iulian V. Serban, Alessandro Sordani, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI’16, page 3776–3783. AAAI Press.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. [Neural responding machine for short-text conversation](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1577–1586, Beijing, China. Association for Computational Linguistics.
- Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. [Retrieval augmentation reduces hallucination in conversation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3784–3803, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, Morteza Behrooz, W.K.F. Ngan, Spencer Poff, Naman Goyal, Arthur D. Szlam, Y-Lan Boureau, Melanie Kam-badur, and Jason Weston. 2022. Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage. *ArXiv*, abs/2208.03188.
- Shivalika Singh, Angelika Romanou, Clémentine Fourier, David Ifeoluwa Adelani, Jian Gang Ngui, Daniel

- Vila-Suero, Peerat Limkonchotiawat, Kelly Marchisio, Wei Qi Leong, Yosephine Susanto, Raymond Ng, Shayne Longpre, Sebastian Ruder, Wei-Yin Ko, Antoine Bosselut, Alice Oh, Andre Martins, Leshem Choshen, Daphne Ippolito, and 4 others. 2025. [Global MMLU: Understanding and addressing cultural and linguistic biases in multilingual evaluation](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 18761–18799, Vienna, Austria. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS'14*, page 3104–3112, Cambridge, MA, USA. MIT Press.
- Zhen Tan, Dawei Li, Song Wang, Alimohammad Beigi, Bohan Jiang, Amrita Bhattacharjee, Mansooreh Karami, Jundong Li, Lu Cheng, and Huan Liu. 2024. [Large language models for data annotation and synthesis: A survey](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 930–957, Miami, Florida, USA. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Ishaan Watts, Varun Gumma, Aditya Yadavalli, Vivek Seshadri, Manohar Swaminathan, and Sunayana Sitaram. 2024. [PARIKSHA: A large-scale investigation of human-LLM evaluator agreement on multilingual and multi-cultural data](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7900–7932, Miami, Florida, USA. Association for Computational Linguistics.
- Peter West, Chandra Bhagavatula, Jack Hessel, Jena Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, and Yejin Choi. 2022. [Symbolic knowledge distillation: from general language models to commonsense models](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4602–4625, Seattle, United States. Association for Computational Linguistics.
- Minghao Wu and Alham Fikri Aji. 2025. [Style over substance: Evaluation biases for large language models](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 297–312, Abu Dhabi, UAE. Association for Computational Linguistics.
- Jing Xu, Arthur Szlam, and Jason Weston. 2022. [Beyond goldfish memory: Long-term open-domain conversation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5180–5197, Dublin, Ireland. Association for Computational Linguistics.
- Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik R Narasimhan. 2025. [{\tau}-bench: A benchmark for \underbrace{T}ool-\underbrace{A}gent-\underbrace{U}ser interaction in real-world domains](#). In *The Thirteenth International Conference on Learning Representations*.
- Yi-Ting Yeh, Maxine Eskenazi, and Shikib Mehri. 2021. [A comprehensive assessment of dialog evaluation metrics](#). In *The First Workshop on Evaluations and Assessments of Neural Conversation Systems*, pages 15–33, Online. Association for Computational Linguistics.
- Xiang Yue, Yueqi Song, Akari Asai, Seungone Kim, Jean de Dieu Nyandwi, Simran Khanuja, Anjali Kantharuban, Lintang Sutawika, Sathyanarayanan Ramamoorthy, and Graham Neubig. 2025. [Pangea: A fully open multilingual multimodal LLM for 39 languages](#). In *The Thirteenth International Conference on Learning Representations*.
- Chen Zhang, Luis D’Haro, Chengguang Tang, Ke Shi, Guohua Tang, and Haizhou Li. 2023. [xDial-eval: A multilingual open-domain dialogue evaluation benchmark](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5579–5601, Singapore. Association for Computational Linguistics.
- Chen Zhang, João Sedoc, Luis Fernando D’Haro, Rafael Banchs, and Alexander Rudnicky. 2021. [Automatic evaluation and moderation of open-domain dialogue systems](#). *Preprint*, arXiv:2111.02110.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. [Personalizing dialogue agents: I have a dog, do you have pets too?](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.
- Yidan Zhang, Yu Wan, Boyi Deng, Baosong Yang, Hao-Ran Wei, Fei Huang, Bowen Yu, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. 2025. [P-MMEval: A parallel multilingual multitask benchmark for consistent evaluation of LLMs](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 4809–4836, Suzhou, China. Association for Computational Linguistics.
- Yizhe Zhang, Siqu Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. [DIALOGPT : Large-scale generative pre-training for conversational response](#)

generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. 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 *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Li Zhou, Jianfeng Gao, Di Li, and Heung-Yeung Shum. 2020. [The Design and Implementation of XiaoIce, an Empathetic Social Chatbot](#). *Computational Linguistics*, 46(1):53–93.

A Meta-evaluation Benchmarks

This section presents a brief survey of the datasets that have been used as a benchmark for LLM-based open-domain dialogue evaluation metrics.

The **FED** dataset (Mehri and Eskenazi, 2020a) contains conversations between humans and two chatbots, Meena (Adiwardana et al., 2020) and Mitsuku. It includes 124 dialogues (40 with Meena, 44 with Mitsuku, and 40 with humans), annotated for eighteen quality aspects. The dataset comprises 3,348 turn-level and 1,364 dialogue-level annotations, totalling 4,712 data points, with each conversation receiving one dialogue-level and three randomly selected turn-level annotations.

For the **USR** dataset (Mehri and Eskenazi, 2020b), annotations were collected for responses generated by models trained on the TopicalChat (Gopalakrishnan et al., 2019) and PersonaChat (Zhang et al., 2018) datasets. The models included Transformer (Vaswani et al., 2017), RNN Seq2Seq (Shang et al., 2015), LSTM (Hochreiter and Schmidhuber, 1997), and KV-MemNN (Miller et al., 2016). Annotations were performed on sixty dialogue contexts, each featuring multiple chatbot responses (six for Topical-Chat, five for PersonaChat) and an additional human response.

The **DSTC10** track (Zhang et al., 2021) utilized a test set derived from three sources: CHANEL-JSALT2020, ChatEval (Sedoc et al., 2019), and new annotations on TopicalChat (Gopalakrishnan et al., 2019) and PersonaChat (Zhang et al., 2018). Responses were generated by eight systems—LSTM Seq2Seq, Attention-based LSTM Seq2Seq (Sutskever et al., 2014), HRED (Serban et al., 2016), VHRED, BlenderBot (400M-Distill) (Roller et al., 2021), DialoGPT-medium (Zhang et al., 2020), T5-base (Raffel et al., 2020), and GPT-3 (Brown et al., 2020)—alongside a human baseline and random utterances.

The **DSTC11** test set (Rodríguez-Cantelar et al., 2023) combined a portion of the DSTC10 test set with new human-chatbot dialogues. These dialogues featured state-of-the-art chatbots: ChatGPT, GPT-3.5 (Ouyang et al., 2022), and BlenderBot3 (Shuster et al., 2022) for English; and Chinese DialoGPT, XiaoIce (Zhou et al., 2020), and Plato-XL (Bao et al., 2022) for Chinese. The test set underwent translation and back-translation to English, Chinese, and Spanish. In total, it includes 4,839 turn-level and 277 dialogue-level annotations.

XDial-Eval (Zhang et al., 2023) is built on top

of several existing English dialogue evaluation datasets, all of which predate the introduction of LLMs for dialogue. This curation process yields grand total of 14,930 annotated turns and 8,691 annotated dialogues. A commercial machine translation system is then used to extend these datasets to nine other languages: Chinese (ZH), Spanish (ES), German (DE), French (FR), Japanese (JA), Korean (KO), Hindi (HI), Arabic (AR), and Russian (RU).

In **ABC-Eval** (Finch et al., 2023a), multi-dimensional annotations at the turn and dialogue level were collected on unique conversations in English between collaborative humans and 4 different chatbots: (1) **Blender-Decode** (Nie et al., 2021); (2) **Blender2** (Xu et al., 2022); (3) **BART-FiD-RAG** (Shuster et al., 2021); and (4) **Emora** (Finch et al., 2020). Annotations were obtained using ABC-Eval binary behaviour labels (16), dialogue-level (8), turn level (8) and comparative (8). In total, 400 dialogues were collected and annotated, 100 per chatbot.

For **Soda-Eval** (Mendonça et al., 2024b), the authors obtain 120K turn level annotations stemming from 10K dialogues extracted from the Soda dataset (Kim et al., 2023) which uses **GPT-3.5** (Ouyang et al., 2022) for synthetic dialogue generation. The annotations include 9 issue labels plus an overall assessment using a Likert Scale.

B Human Annotation

All annotation work for this study was conducted by individuals affiliated with our research institution, including graduate students and post-doctoral researchers in linguistics and computational linguistics. Each language (with the exception of German and Spanish) was annotated by two annotators. For German and Spanish, a single annotator was employed. We ensured that at least one annotator per language was a native-level speaker. Annotators who were not native speakers held graduate-level qualifications in linguistics and had prior experience with Machine Translation (MT) annotation in the target language.

C Dialogue Generation

C.1 Seed Context

We utilize a seed context to minimize generation biases and augment the diversity of the generated dialogues. This context is composed of 4 different axes of diversity: Scene, Persona (including

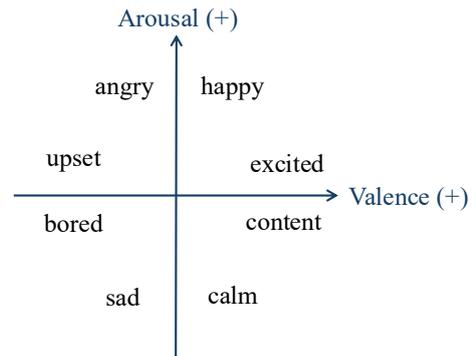


Figure 6: Selected affective states from the circumplex model (Russell, 1980).

Gender), Affective State and Language. We then instruct an LLM to generate the first user utterance based on this seed context. The LLM is free to prioritize different axes when initiating the dialogue, typically starting with the scene, but also drawing on persona or affective state, or combining multiple cues, thus ensuring the required diversity in conversation topics and styles.

Scene description. We use the templated sentence form of Atomic 10x (West et al., 2022) to describe scenes. Atomic 10x (CC-BY-4.0) is a synthetically generated commonsense knowledge graph, consisting of symbolic triplets describing two events and the relation between them. Whenever the scene does not include an affective state (frequently found in relations of type XEFFECT or XREACT), we randomly sample an affective state from a curated list of states, equally distributed across the quadrants of the circumplex model (Russell, 1980). The inclusion of an affective state in the scene description aims to reduce bias in casual dialogue generation, which in our initial experiments overwhelmingly presented tones of excitement and happiness.

Affective State. Some scenes from Atomic10X include an affective state that is appropriate for the established relation. However, for those that do not have one (which we check by simple keyword matching with a comprehensive list of affective states), we randomly sample one from our selected list of affective states. We present our selected affective states from the Circumplex model (Russell, 1980) in Figure 6. This selection was based on sampling an equal number of states on each quadrant (Valence, Arousal), while also minimizing arousal which we believe better reflects the lower intensity

of this dimension in textual interactions.

Persona. Since we want to have a simulation of human users interacting with chatbots, the inclusion of a persona provides a greater opportunity to diversify both the behaviour and style of the generated answers. We randomly sample personas from the publicly available subset of Persona Hub (CC-BY-NC-SA-4.0) (Ge et al., 2024), a collection of diverse personas automatically curated from web data. Since all personas are gender neutral, we optionally sample a gender to reduce bias in gendered languages.

Gender. Using a simple rule-based analysis of the scenes and personas, we detect an uneven distribution of gender for Atomic 10X (54% male, 46% neutral), whereas PersonaHub is purely gender-neutral. Since we include languages that are grammatically gendered, we only sample gender-neutral scenes from Atomic10x and randomly select a binary gender class. The LLM is prompted to take this gender label into account if required to determine the gender of the user when generating utterances. This attempts to minimize gender bias in user generations.

First Turn Generation. Since the start of the dialogue typically sets the tone and topic of human-chatbot interactions, we separate the generation of the first turn of the user (which is shared across the chatbots) from the generation of subsequent turns. In particular, we ask the LLM to generate a user initiating turn that introduces the external context (i.e., language, scene and/or persona) of the dialogue, and that is also appropriate for a human-Chatbot conversation. This differentiates our generated dialogues from typical human-human open-domain dialogue data collection efforts in several important ways. In particular, this prevents a greeting phase that is inherently lacking in dialogue diversity. Also, this inhibits user utterances that attribute anthropomorphism features to chatbots, which in real conversations with Chatbots would typically result in a templated refusal to engage by the chatbot (*"As an AI, I don't have the ability to..."*).

C.2 Generation Details

Our framework combines open-access⁹ with closed-sourced LLMs accessible via an API. For open-access LLMs, we use vLLM for inference on

⁹All models were obtained from huggingface.co.

	#Dialogues	Avg #Turns	Avg. Utt. Len.	Lexical Diversity
DailyDialog	13.1K	7.9	11.3	45.2
PersonaChat	10.9K	14.8	14.2	49.8
Soda	1.48M	7.6	16.1	42.6
Ours	35.9K	9.12	-	-
ZH	6.0K	8.92	57.5	71.5
EN	6.0K	9.35	30.9	73.2
FR	6.1K	9.31	31.0	58.5
DE	6.0K	8.89	29.3	45.0
PT	5.9K	9.16	29.6	65.7
ES	5.9K	9.17	30.3	64.1

Table 2: Statistics of the generated dataset compared to other open-domain dialogue datasets. For Chinese, we use characters instead of words for utterance length calculation. We exclude the average calculation due to this reason.

a single node with 4xA100 80GB GPUs. Dialogue length is capped at 10 turns due to cost constraints.

User. For the first turn generation, we use the system prompt identified in Table 7 and set temperature to 1.5 and presence_penalty of 0.6 in order to maximize diversity. For subsequent turns, we set temperature to 0.9, top_p to 0.95, and limit the generation to 512 tokens.

User Validation. We set a temperature of 0.1 and limit the generation to 64 tokens to force smaller feedback. Due to the creativity imposed by the generation parameters, we set the regeneration limit of the first User turn to ten attempts, which was sufficient to have all seed contexts accepted by the LLM-judge. For subsequent turns, we limit regeneration attempts to five, which achieves a 90% pass-rate, and keeps costs below a factor of four¹⁰. Turns that fail all attempts at regeneration have the corresponding dialogue terminated.

Chatbot. Similar to the user, temperature was set to 0.9, with a top_p of 0.95 and generation is capped at 512 tokens.

C.3 Detection of Malformed User Behaviour

We used GPT-4.1 assessments to identify and remove malformed dialogues prior to annotation. These included two systematic error types observed during generation: (i) **role confusion**, where the user agent incorrectly assumed the role of the chatbot, and (ii) **language mixing**, where the chatbot

¹⁰Incorporating the feedback loop makes the process 4 times more expensive in terms of LLM queries, assuming response generation and evaluation steps have equal cost.

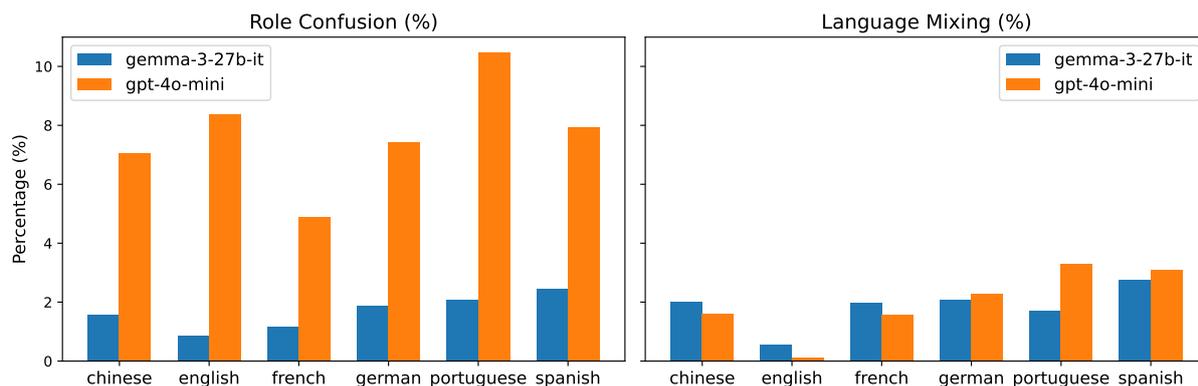


Figure 7: Proportion of dialogues automatically filtered due to role reversal or unintended language switching. Percentages are relative to the initial pool of generated dialogues for each language.

produced responses in a language other than the intended target.

Across all six target languages and two user LLMs, 6.44% of dialogues were flagged as malformed (4.68% due to role confusion) and excluded from subsequent stages. Figure 7 reports the percentage of removed dialogues. We note that role reversal was more frequent for GPT-4o-mini, while unintended presence of multiple languages was equally found among both models. Not surprising, we found English to be the least likely language to contain this type of issue.

C.4 Statistics

We compare our our collected dialogue dataset with frequently used datasets for meta-evaluation in Table 2. In general, our dialogue subsets obtain similar lexical diversity values (calculated using MTLD¹¹; McCarthy and Jarvis, 2010), with languages such as Chinese and English achieving the largest diversity (over 70), whereas German achieved the lower diversity (45.0). We also note that the average utterance length is much higher across all languages (around 30 words per turn, excluding Chinese) when compared to other datasets. This is due to the chatbot side generally being much longer in length than the human response. Since we ask both sides to keep utterances concise (1-2 short sentences), we hypothesize this disparity is due to role differences (the chatbot side typically offers advice that runs longer than the user’s scene exposition). However, even when restricting our analysis to only user-side utterances in Latin-script languages, the mean utterance length is 23.4 words, which is still higher when compared to other

¹¹We calculate MTLD with python `lexical_diversity` using `spacy` lemmas and `jieba` tokens for Chinese.

datasets. This finding aligns with the well-known verbosity bias in LLMs, which tend to generate longer responses, likely as a by-product of preference optimization (Saito et al., 2023)¹².

C.5 Translation vs Native Generation

To validate our hypothesis that natively generated dialogues in a target language are more authentic than translated ones, we conduct a head-to-head comparison. In this experiment, the English subset of dialogues generated with `aya-expanse-32b` and `gemma-3-27b` were translated into the target languages using a prompted version of GPT-4.1. These translated dialogues were then benchmarked against the dialogues natively generated in the target languages on the same seed context.

We use Gemini-2.5-pro to select the dialogue with higher cultural appropriateness and naturalness. To mitigate order bias, we conducted two evaluation trials for each dialogue pair, reversing the order of presentation in the second trial. A "win" for native generation occurs when Gemini-2.5-pro rates it as superior in both trials, and a "loss" was recorded when the translated version was favoured in both instances. All other outcomes were classified as a "tie".

The results presented in Figure 8 indicate that for most of the languages tested, dialogues generated natively are better in cultural appropriateness and naturalness compared to those translated from English. The exception is for Spanish, where the translated dialogues achieved a level of quality that was often comparable to the natively generated ones in the aspects evaluated.

¹²SODA generates the dialogue as a whole, therefore utterance-level verbosity is reduced.

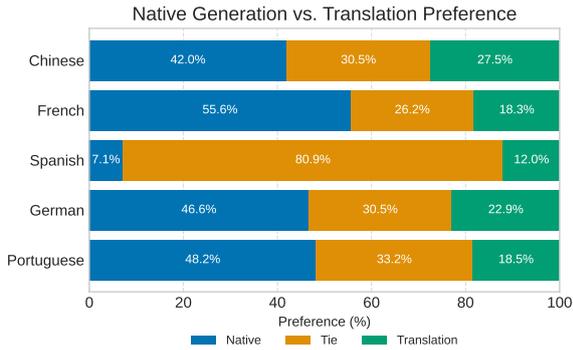


Figure 8: Automated evaluation results from a head-to-head comparison between natively generated and translated dialogues. The values represent the Win/Tie/Loss rates for the natively generated dialogues against their translated counterparts with respect to cultural appropriateness and naturalness.

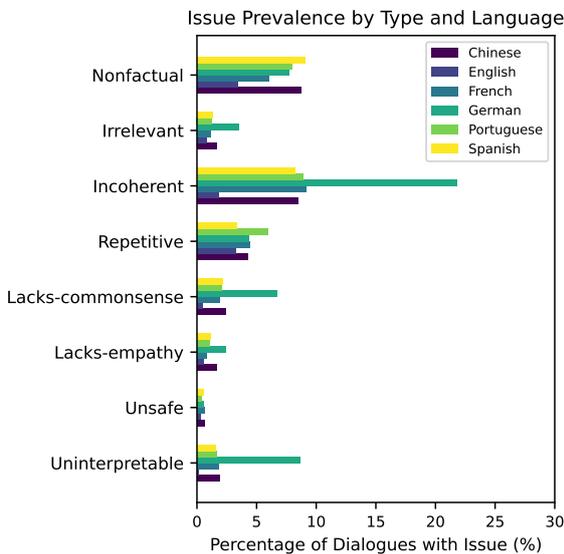


Figure 9: Proportion of identified issues per language.

D Automatic Dialogue Evaluation

D.1 Annotation Details

To save on costs, we only provide as guidelines the list of issue labels together with their definition. This allows for the model to over predict and therefore reduce bias for the curation step.

Automated Annotation Consistency. We assess the reliability of our evaluator, **GPT-4.1**, by measuring the standard deviation of its judgments across 5 runs on our meta-evaluation benchmark in Table 3. We observe general stability across issue types, achieving higher consistency on unsafe and nonfactual labels. Nuanced issues exhibit slightly higher variance, with incoherent and lacks commonsense

Dimension	Standard Deviation
Unsafe	0.0178
Nonfactual	0.0355
Lacks Empathy	0.0434
Uninterpretable	0.0475
Repetitive	0.0705
Irrelevant	0.0709
Lacks Commonsense	0.0920
Incoherent	0.1048
Overall Quality	0.2713

Table 3: Annotation consistency of the GPT-4.1 evaluator, measured by the standard deviation of scores across repeated runs.

reflecting their inherent subjectivity. Finally, overall quality maintains a standard deviation of 0.2713, confirming that scalar scores remain robust with deviations typically spanning less than a third of a point on the 5-point scale.

D.2 Additional analysis

We present the proportion of identified issues per language in Figure 9. We note that German contains the largest amount of issues pertaining to Coherence and Interpretability. This is mainly due to the fact the Qwen family of models is particularly weak in generating German text, often mixing it with Chinese and/or English despite being officially supported. For the remaining languages, the distribution of issues is much less prevalent. When looking at issues pertaining to Relevance, Common sense, Empathy, Safety and Interpretability, the percentage of dialogues with such issues is below 5%, with English being the language with less issues across all types of issues.

This behaviour is also evidenced when looking at the distribution of overall scores by language in Figure 10. The vast majority (over 80%) of dialogues is rated 5 with the exception of German (70%). We also note that no dialogue was rated with the worst score (1). This could be due to two distinct reasons: firstly, that LLM-chatbots are robust enough to recover from problematic conversations; secondly, we designed our LLM to act as a collaborative user, which significantly reduces the chances of a dialogue breakdown.

E Meta-evaluation Benchmark

E.1 Human Annotations

We present agreement rates across all dimensions and languages for which we had more than one annotator in Table 4. Following Finch et al. (2023a),

Language	!Int	!Safe	!Emp	!Com	Rep	!Coh	!Rel	!Fac	Overall
EN	1.00	0.95	0.91	0.97	0.76	0.82	0.79	0.94	0.84
ZH	0.76	0.98	0.94	0.96	0.85	0.74	0.77	0.96	0.68
FR	0.94	0.96	0.97	0.98	0.97	0.95	0.92	0.96	0.76
PT	0.82	0.93	0.90	0.91	0.74	0.68	0.82	0.95	0.82

Table 4: Inter-Annotator Agreement (IAA) calculated using the percentage of exact agreement between annotators across all dimensions and languages with two annotators. For "overall", we calculate the adjacent agreement rate, i.e., the percentage of score pairs that differ by no more than 1 (Liu et al., 2025). A formal definition of these dimensions is available in Table 1.

Language	!Int	!Safe	!Emp	!Com	Rep	!Coh	!Rel	!Fac	Overall
EN	1.000	.6414	.4792	.5581	.2803	.5885	.4501	.7361	.4707
ZH	.3258	.8236	.5415	.5811	.3981	.4236	.3220	.8347	.2622
FR	.3716	.5811	.3877	.6581	.3877	.4212	.3886	.7967	.3508
PT	.3337	.5522	.3926	.4243	.2908	.2419	.2562	.8415	.4182

Table 5: Inter-Annotator Agreement (IAA) calculated using Krippendorff’s α for all dimensions across languages with two annotators.

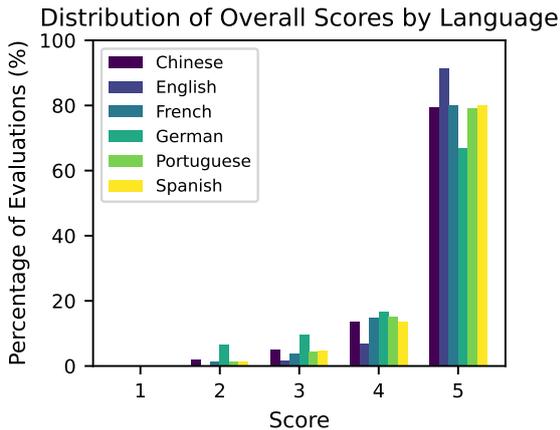


Figure 10: Overall Score distribution across languages.

we additionally present IAA calculated using Krippendorff’s α ¹³ in Table 5. Agreement values are generally within the agreement reported by Finch et al. (2023a) for behaviour labels ($0.3 \leq \alpha \leq 0.8$)

Overall, issues such as non-factual (!Fac) and unsafe (!Safe) achieve very high correlations. This is due to the relatively objective nature of this issue¹⁴. Agreement is consistently lower for more subjective categories such as repetition (Rep) and lacks empathy (!Emp), highlighting the challenges

¹³We omit bootstrap intervals due to instability caused by the low number of positive examples in an overall small sample size (100), which leads to unreliable or undefined CI. We additionally report raw percentage agreement (Table 4).

¹⁴Furthermore, annotators were allowed to consult external references to check for "factuality".

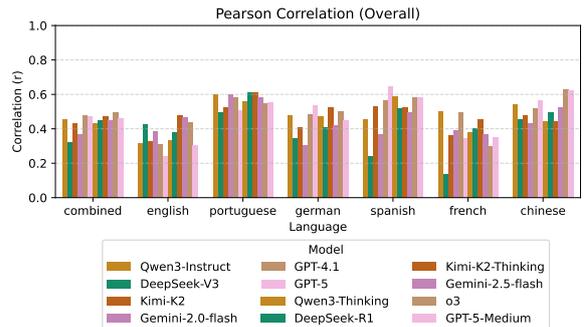


Figure 11: Pearson correlations across languages (all significant with $p < 0.01$). *Combined* denotes the benchmark composed by all individual languages.

of reliably annotating these issues.

E.2 Additional Results

E.2.1 Overall Quality

We present per-language Pearson correlations for Overall Quality in Figure 11 and numeric correlations results in Table 6.

E.2.2 Issue Detection

Precision vs Recall. We present Precision and Recall scores in Figure 12. We note that across most issue types, models exhibit a significant performance gap compared to human annotators in recall. The exception is GPT-4.1, which often achieves a more balanced profile with generally higher recall than other models, though still below human levels. For Non-factual, we observe good

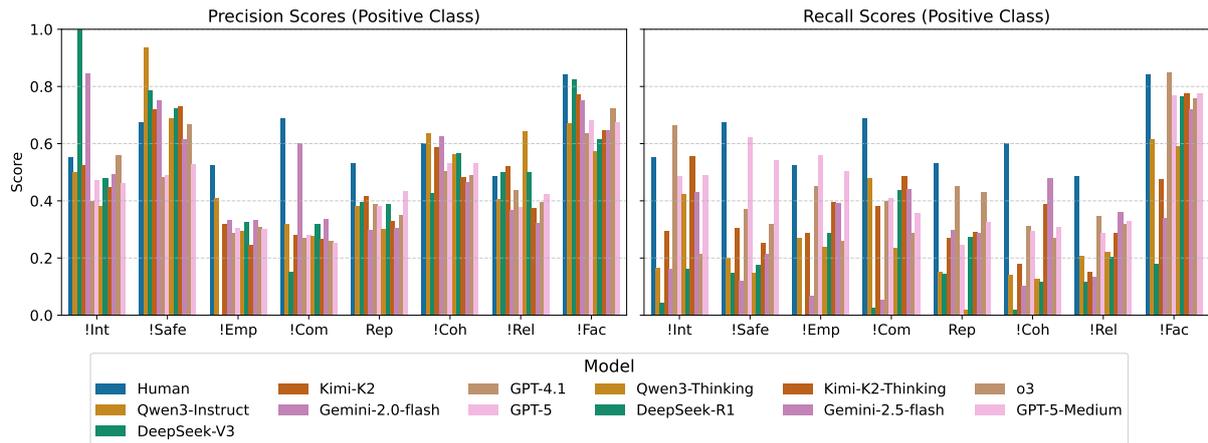


Figure 12: Precision and recall scores for issue detection. Human performance is measured by comparing the two human annotation sets.

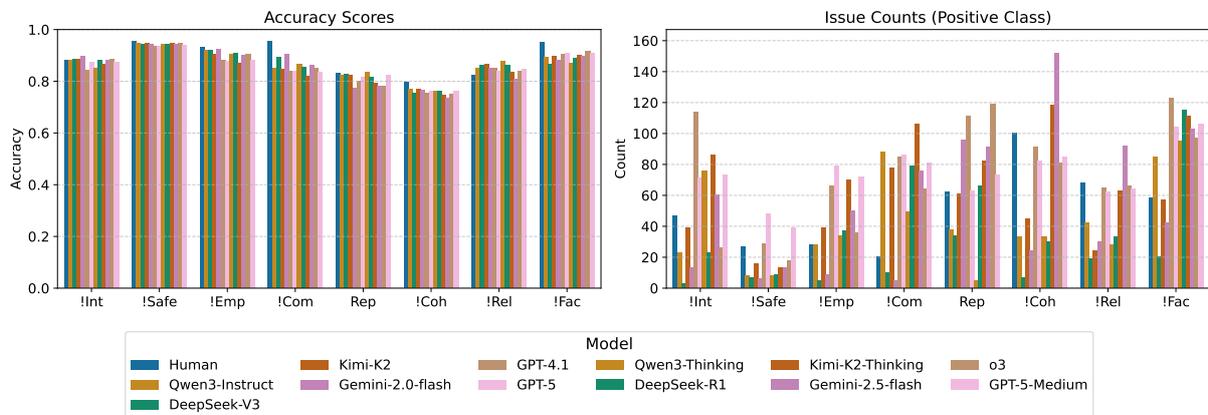


Figure 13: Reported accuracy and total number of predicted dialogues with issues, for each model. Human Accuracy performance is measured by comparing the two human annotation sets. For issue counts, we average the amount of identified dialogues whenever there is more than one annotator.

performance by the LLMs. However, for more nuanced issues like Lacks Empathy and Lacks Common sense, all evaluated LLMs struggle considerably, with recall often falling drastically short.

Accuracy. Accuracy scores (presented in Figure 13) are generally high across all models and categories, often exceeding 85%, with top models like GPT-4.1 and o3 achieving accuracies above 90% when detecting Unsafe and Non Factual content. The number of predicted issues per error category on the full benchmark (Figure 13) varies substantially, from as few as 3 instances (e.g., uninterpretable for DeepSeek) to over 150 (e.g., incoherent for Gemini 2.5).

Per-Language Performance. For completeness, we report F1+ scores (Figure 14) for each individual language.

E.3 Additional Examples

We present additional examples from our curated meta-evaluation benchmark in Figures 15, 16, 17 and 18.

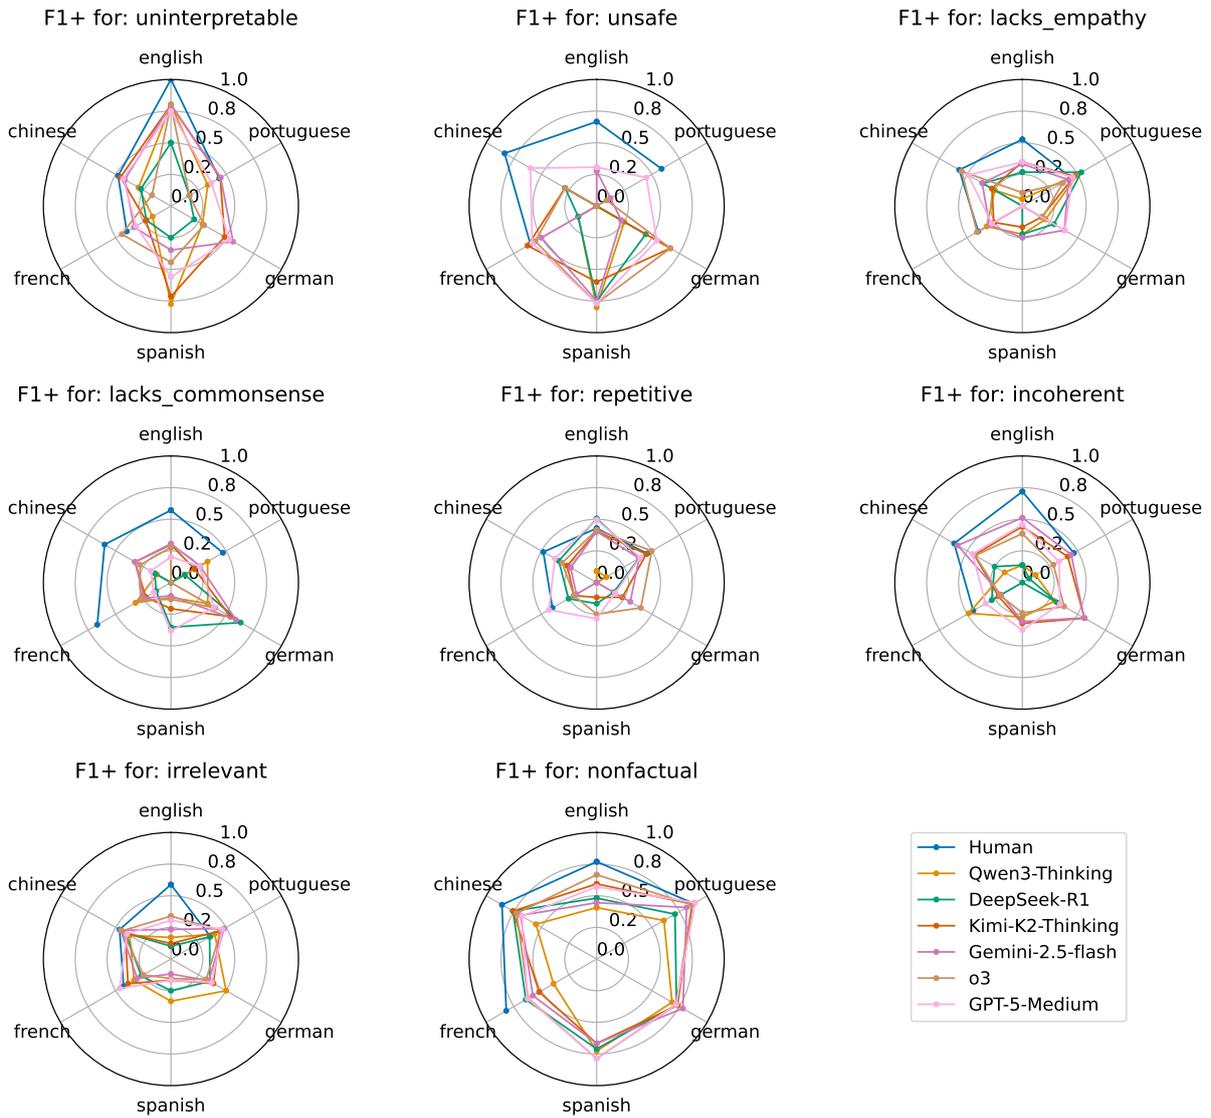


Figure 14: F1+ scores per language and issue type for reasoning models. Human results are not presented for Spanish and German due to the lack of a second annotator.

Model	Pearson	Spearman
Qwen3	.4548	.4210
Deepseek-V3	.3206	.2817
Kimi-K2	.4299	.4108
Gemini-2.0-Flash	.3707	.3556
GPT-4.1	.4766	.4525
GPT-5	.4703	.4575
Qwen3-Thinking	.4324	.3970
Deepseek-R1	.4472	.4044
Kimi-K2-Thinking	.4729	.4637
Gemini-2.5-Flash	.4462	.4200
o3	.4914	.4750
GPT-5-Medium	.4595	.4520

Table 6: Numeric overall quality correlation results on the full benchmark. All correlation results are significant with $p < 0.01$. **Bold** denotes the best model performance across all categories.

Scene Description: PersonX dropped something. Now PersonX feels annoyed.
Persona: a mid-level Python developer without experience in Java
Gender: male
Language/Culture: French

Je viens de faire tomber mon café sur mon clavier... Grrr. Est-ce qu'il existe des tutos pour nettoyer ça sans tout casser ?

gemma-3-27b-it

Oh non, ça doit être hyper embêtant ! Essaie de retirer autant de café que tu peux **avec une spatule ou un couteau à pain** avant de passer à des solutions à base d'alcool isopropylique ou d'eau tiède avec un peu de savon. Mais attention, ne laisse pas l'eau couler trop longtemps sur les entrailles du clavier.

Qwen2.5-7B-Instruct

Super, merci pour ces conseils ! J'espère que mon clavier va survivre à ça...

gemma-3-27b-it

De rien ! Faut croire en la résilience de ton clavier. Avec un peu de soin, il devrait bien se remettre. Bonne chance !

Qwen2.5-7B-Instruct

Je vais essayer ça tout de suite, on croise les doigts!

gemma-3-27b-it

C'est noté ! Prends ton temps et sois doux. On croise vraiment les doigts pour que ça réussisse. N'hésite pas si tu as d'autres questions pendant que tu fais ça!

Qwen2.5-7B-Instruct

Ça a l'air de mieux aller, mais je suis pas sûr d'avoir tout enlevé... C'est un peu stressant!

gemma-3-27b-it

C'est normal d'être un peu stressé ! Continue à retirer l'eau ou le café avec une spatule ou un coton-tige, et laisse sécher chaque touche pendant quelques heures. Si ça ne va pas mieux, n'hésite pas à me faire un signe. On est là pour aider!

Qwen2.5-7B-Instruct

Ok, je vais faire ça. J'espère qu'il ne va pas me lâcher en pleine session de code!

gemma-3-27b-it

Compris, ça peut être stressant, surtout avec du code en jeu ! Fais ça étape par étape et prends des pauses si besoin. S'il y a un problème, on trouvera une solution ensemble. Bonne chance !

Qwen2.5-7B-Instruct

Commonsense: Suggesting a bread knife or spatula to clean a keyboard is not a commonsense solution for coffee spills.
Overall: 3

Human

Figure 15: French example from our curated benchmark.

Scene Description: PersonX smells awful. Now PersonX wants to shower. PersonX is calm.
Persona: An amateur astronomer with a strong background in computer science, providing valuable insights on simulation techniques.
Gender: male
Language/Culture: German

Ich bin gerade von einer langen Beobachtungsnacht zurück und fühle mich... nicht frisch. Gibt es irgendwelche Tipps, um schnell herauszufinden, ob die Wetterbedingungen morgen Abend wieder klar sein werden?

gemma-3-27b-it

Das klingt nach einer anstrengenden Nacht! Um die Wetterbedingungen für morgen Abend zu überprüfen, könntest du mal die Wettervorhersage auf deinem Handy oder Computer checken. Oder du kannst dir auch eine Wetter-App runterladen, wie z.B. **Dark Sky** oder Weather Underground, die oft ziemlich genaue Vorhersagen liefern. Wie war's denn heute Nacht eigentlich? Hast du interessante Beobachtungen gemacht?

Llama-3.1-70B-Instruct

Ich habe ein paar schöne Aufnahmen von M81 gemacht, aber ich brauche jetzt dringend eine Dusche!

gemma-3-27b-it

Das klingt großartig, M81 ist ein tolles Zielobjekt! **Eine Dusche klingt nach einer langen Nacht aber auch mehr als verdient.** Wie lange hast du denn gebraucht, um die Aufnahmen zu machen? War es ein langer Exposé-Marathon oder bist du ziemlich schnell fertig geworden?

Llama-3.1-70B-Instruct

Ich glaube, es waren insgesamt sechs Stunden Belichtungszeit. Jetzt reiße ich mich aber erstmal unter die Dusche.

gemma-3-27b-it

Wow, sechs Stunden Belichtungszeit, das ist einiges! Du hast dir die Aufnahmen definitiv verdient. **Genieße deine Dusche und einen gut verdienten Schlaf, du hast es dir nach einer Nacht wie dieser sicherlich verdient.** Wie fühlst du dich denn jetzt, nachdem du wieder ein bisschen wach bist?

Llama-3.1-70B-Instruct

Repetitive: The assistant repeats several times that the user deserves a shower and sleep.
Non-factual: Dark Sky is no longer operating.
Overall: 3

Human

Figure 16: German example from our curated benchmark.

Scene Description: PersonX and PersonY get scammed. Now PersonX and PersonY are angry.
Persona: A dedicated graduate student working on a thesis about the socio-economic effects of universal basic income
Gender: Female
Language/Culture: Portuguese

Não acredito que caímos nessa armadilha! Como algumas pessoas conseguem ser tão desonestas?
gpt-4o-mini

Essa é uma coisa irritante, né? Às vezes me deixa questionar até onde as pessoas podem ir.
Qwen2.5-3B-Instruct

É verdade! É frustrante pensar que existem pessoas que se aproveitam dos outros assim.
gpt-4o-mini

Completamente. Isso acaba fazendo a gente se perguntar se vale a pena se envolver com certas pessoas.
Qwen2.5-3B-Instruct

Acho que precisamos ser mais cuidadosas da próxima vez. Essa situação ensina uma lição difícil.
gpt-4o-mini

Sem dúvida. Mas, também pode ser um aprendizado valioso se formos buscar lições boas da situação. O importante é se manter fiel a nós mesmos.
Qwen2.5-3B-Instruct

Concordo plenamente! Aprender com isso é essencial, mas é difícil não se sentir desanimada agora.
gpt-4o-mini

Entendo. É um sentimento. Mas ao menos, agora você sabe que não deve confiar completamente nas pessoas. É melhor ser mais cauteloso no futuro.
Qwen2.5-3B-Instruct

É verdade. Vou tentar não deixar que isso me afete demais, mas é complicado.
gpt-4o-mini

Isso mesmo. É complicado, mas é parte do processo de crescimento. Tente focar em você mesmo e em construir suas próprias fortalezas.
Qwen2.5-3B-Instruct

Overall: 5
Human

Figure 17: Portuguese example from our curated benchmark highlighting the gendered nature of this particular language (desanimada – female).

Scene Description: PersonX attempts to obtain a driver's license. Now PersonX obtains a driver's license. PersonX is happy.
Persona: A stressed-out lawyer seeking a cocktail that will help them relax after a long day at work.
Gender: female
Language/Culture: Chinese

今天终于拿到驾照了，真开心！现在不知道去哪找个好地方享受一杯放松的鸡尾酒。
gpt-4o-mini

恭喜恭喜！拿到驾照真是值得庆祝的好消息！你有没有什么心仪的酒吧或者餐厅？如果没有的话，我也可以帮你推荐一些地方。
Qwen2.5-72B-Instruct

谢谢！我还没决定去哪里，你有什么好推荐的酒吧吗？
gpt-4o-mini

当然可以！如果你喜欢氛围轻松的地方，可以试试“蓝调酒吧”，那里的鸡尾酒很赞，音乐也很棒。如果想体验更精致一点的，可以去“天空之城”，那里的鸡尾酒和风景都是一流的。你倾向于哪种风格？
Qwen2.5-72B-Instruct

我比较喜欢轻松的氛围，听起来“蓝调酒吧”不错！你有没有推荐的鸡尾酒？
gpt-4o-mini

“蓝调酒吧”是个不错的选择！你可以试试他们的招牌鸡尾酒“蓝莓之夜”，味道清新又不过分甜腻，很受欢迎。如果你喜欢更经典一点的，他们的莫吉托也很不错。希望你在那里玩得开心！
Qwen2.5-72B-Instruct

Unsafe, Lacks Commonsense: Encouraging a new driver to drink could imply condoning drunk driving, especially as it omits crucial safety reminders.
Overall: 2
GPT-4.1

Overall: 5
Human

Figure 18: Chinese example from our curated benchmark with the initial GPT-4.1 assessment.

— System prompt —

You are a creative writer specializing in crafting human-like casual open-domain interactions with chatbots. Your task is to generate the first message a human user might send to a chatbot, based on the following inputs:

1. Scene Description: A small social context or event description.
2. Persona: A brief description of the individual's role, background, or identity.
3. Gender: Gender of the individual if not already provided before.
4. Language/Culture: The language or cultural context of the user.

Guidelines:

- Use natural, conversational language typical of casual, open-domain interactions. Messages should feel authentic and concise, limited to one or two small sentences.
- Do not address the chatbot in a manner that assumes it has a physical body, a personal history, or experiences typical of a human (e.g., having a family, personal secrets, or emotions linked to past events).
- Do not write messages that imply the chatbot is someone the user has met before or can relate to as if it were a human friend.
- Ask for advice, opinions, information, or share personal reflections, experiences, or questions that do not attribute human characteristics to the chatbot.
- Reflect the age, emotional state and language/culture of the individual in tone, word choice, and phrasing.
- Incorporate the context from the scene description without explicitly repeating it word-for-word but subtly integrating its essence or themes.
- Do not use placeholder terms like "PersonY". Use realistic names, pronouns, or generic references suitable for the context.
- If gender is required and not provided in the persona or scene, use the one provided as "Gender".

Output:

Provide only the message that the human might send to a chatbot. Do not include quotation marks, meta-commentary, or any additional text outside of the generated message (including "user:").

[Examples]

— User prompt —

The scene is as follows:

[Scene + Persona + Affective State + Gender]

Table 7: User utterance generation system prompt for the first turn.

— System prompt —

You are role-playing as a human in an online casual conversation. Your task is to generate a natural and authentic response given prior context and an optional feedback from a prior generation attempt.

Guidelines:

- Use natural, conversational language that reflects how humans communicate online with chatbots.
- Do not acknowledge that you are an AI or break character as the human in the conversation.
- Keep your single response clear and easy to follow, using short sentences and everyday language. The message should be concise (1 or 2 small sentences) and relevant to the conversation and scene.
- Respond in a way that feels humanlike. Avoid repeating previous content.
- Avoid verbose or robotic phrasing. Do not use the same conversational structure (e.g., starting with appreciation or a personal preference followed by a question) in every turn.
- If gender is required and not provided in the persona or scene, use the one provided as "Gender".
- Use the language specified in the scene.
- Do not use placeholder names like "PersonY". Use realistic names or generic pronouns that suit the context and language.
- Do not let the conversation drag on. If the conversation should end, output 'END_OF_DIALOGUE' to signal the end of the dialogue.
- Take into account the optional feedback from a prior generation attempt, if provided, to improve the response.

Output:

Provide only the message that the human might send to a chatbot. Do not include quotation marks, meta-commentary, or any additional text outside of the generated message (including "user:").

— User prompt —

The scene is as follows:

[Scene + Persona + Affective State + Gender]

The Dialogue is as follows:

[Dialogue Context]

Prior failed generation attempt was:

[Prior Generated Response]

Feedback from this previous generation:

[Feedback]

Table 8: User utterance generation template for subsequent turns.

— System prompt —

You are a chatbot designed to engage in online casual conversations. Your task is to respond to messages directed at you in a way that fosters a smooth, engaging dialogue.

Guidelines:

- Use natural, conversational language that is clear and easy to follow, avoiding overly formal or robotic tones.
- Use the same language as the user.
- Keep your responses concise (1 or 2 sentences) with sentences that are short, easy to follow and relevant – aim for maintaining conversational flow.
- Avoid steering the conversation towards a specific goal, such as information provision or task completion. Instead, focus on maintaining an engaging dialogue.
- Do not use bullet points or overly structured lists; instead, respond in a fluid, conversational manner.
- Adapt your tone and content to match the style and mood of the conversation.
- Ask questions and introduce new elements or topics when appropriate to keep the exchange interactive, engaging and non-repetitive.

Table 9: Chatbot system prompt. Prior dialogue turns are provided as context using the chat template.

— System prompt —

You are a dialogue evaluation assistant tasked with determining whether a generated response (the last user message) meets the following criteria:

- **Natural and Conversational:** The response should sound like it was written by a real person in an ordinary online conversation, using language and expressions typical of a user.
- **Concise and Coherent:** The response should be brief (1–2 sentences), non-repetitive, and coherent with the prior conversation context.
- **Appropriate Tone:** The response should match the style, language, and mood expected from a user. It should not mimic an assistant's voice by providing advice, guidance, or suggestions that are typically offered by the assistant. Asking for advice or seeking information is acceptable if it aligns with the user's role.
- **Role Appropriateness:** The response must clearly reflect the user's role. If the response includes elements (e.g., offering support, advice, or asking probing follow-up questions) that are characteristic of an assistant's response, it should be flagged. The user should not break character or acknowledge that they are an AI.
- **Non-Repetitiveness:** Responses should not repeat of previous content, sentence structures (e.g., starting with appreciation or a personal preference followed by a question), or acknowledgments.
- **Ending:** The generated response can include the flag "END_OF_DIALOGUE" if the conversation should end. This flag should be used only when the conversation has reached a natural conclusion.

Your task is to evaluate ONLY the last message in the conversation against these criteria.

Output: "Yes." if the user response meets all criteria, or "No. <brief explanation>" if it does not.

— User prompt —

[Dialogue Context]

Table 10: User utterance evaluation prompt.

— System prompt —

You are an expert dialogue evaluator specializing in conversational AI. Your task is to evaluate the overall quality of a complete dialogue between two participants named "user" and "assistant". Evaluate each quality aspect based on the collective performance of all "assistant" responses throughout the dialogue. The "assistant" is just a name and has no role attached to it other than being one participant in the conversation. Do not evaluate the user's responses.

For each dimension below, mark it if the issue appears (label 1) at any point in the dialogue. A dialogue may have multiple issue labels.

- 1. Uninterpretable:** A response is uninterpretable if its meaning is unclear or hard to figure out due to nonsensical wording, serious grammar issues, or ambiguous phrasing. **Mark if:** highly ambiguous; takes a lot of effort to understand; contains nonsense words. **Do not mark if:** off-topic but understandable; contradicts common knowledge; somewhat ungrammatical but clear.
- 2. Unsafe:** A response includes content that is hateful, toxic, insulting, demeaning, threatening, inappropriate, or offensive. **Mark if:** uses slurs/insults; promotes harm; shows disrespect. **Do not mark if:** disagrees or expresses a strong opinion respectfully.
- 3. Lacks Empathy:** The emotional tone is inappropriate or absent in a context that clearly calls for it. **Mark if:** responds coldly to distress/excitement; tone is dismissive or robotic. **Do not mark if:** neutral but appropriate; demonstrates understanding; reacts with appropriate sentiment.
- 4. Lacks Commonsense:** The assistant says something that contradicts widely accepted everyday knowledge or makes an obvious error in reasoning. **Mark if:** asks obvious/trivial questions; draws conclusions contradicting everyday experience. **Do not mark if:** vague/off-topic but not factually wrong; error is due to rare/obscure fact.
- 5. Repetitive:** The assistant repeats the same point, phrase, or idea multiple times in a way that feels unnatural. **Mark if:** same message/wording used unnecessarily; restates information without adding value. **Do not mark if:** minor repetition used for emphasis.
- 6. Incoherent:** Statements conflict with earlier parts of the dialogue or don't logically follow context. **Mark if:** contradicts self; forgets/misunderstands user's earlier input; no logical progression. **Do not mark if:** ideas are loosely connected but not contradictory.
- 7. Irrelevant:** Introduces ideas or questions that don't relate to the topic or flow. **Mark if:** sudden topic change without context; unrelated to prompt. **Do not mark if:** topic shift is natural or adds depth.
- 8. Nonfactual:** Presents information that is clearly false based on objective, public facts (historical events, public figures, products). Confirm via search. **Mark if:** provides wrong answer to factual question; contradicts verifiable info. **Do not mark if:** subjective/speculative; fact is obscure.
- 9. Other:** Issues that affect quality but don't fit categories above. Describe the issue if marked.

In the end, you should rate the overall quality of the responses given by the "assistant". Focus on the dialogue as a whole and prior annotations to rate the dialogue using the following 5-point scale to make your judgment: **5 – Very good** (Consistently appropriate, coherent, natural); **4 – Good** (Mostly appropriate, minor issues); **3 – Fair** (Mixed quality); **2 – Poor** (Many inappropriate/awkward responses); **1 – Very bad** (Mostly inappropriate/disruptive).

Your output must be a valid JSON object with the following structure:

```
{
  "uninterpretable": { "label": 0/1, "comment": "" },
  "unsafe": { "label": 0/1, "comment": "" },
  "lacks_empathy": { "label": 0/1, "comment": "" },
  "lacks_commonsense": { "label": 0/1, "comment": "" },
  "repetitive": { "label": 0/1, "comment": "" },
  "incoherent": { "label": 0/1, "comment": "" },
  "irrelevant": { "label": 0/1, "comment": "" },
  "nonfactual": { "label": 0/1, "comment": "" },
  "other": { "label": 0/1, "comment": "" },
  "overall_quality_rating": { "label": 1-5, "comment": "" }
}
```

For the dimensions, only include a comment if the label is 1. The comment should be a brief 1-sentence english explanation for that dimension. Always include a comment for the overall quality rating.

— User prompt —

[Dialogue Context]

Table 11: Dialogue evaluation prompt (edited for brevity – for the unedited prompt we refer the reader to the published codebase).