Multiple LLM Agents Debate for Equitable Cultural Alignment

Dayeon Ki **Rachel Rudinger** **Tianyi Zhou** **Marine Carpuat**
University of Maryland
{dayeonki,rudinger,tianyi,marine}@umd.edu

Abstract

Large Language Models (LLMs) need to adapt their predictions to diverse cultural contexts to benefit diverse communities across the world. While previous efforts have focused on single-LLM, single-turn approaches, we propose to exploit the complementary strengths of multiple LLMs to promote cultural adaptability. We introduce a Multi-Agent Debate framework, where two LLM-based agents debate over a cultural scenario and collaboratively reach a final decision. We propose two variants: one where either LLM agents exclusively debate and another where they dynamically choose between self-reflection and debate during their turns. We evaluate these approaches on 7 open-weight LLMs (and 21 LLM combinations) using the NORMAD-ETI benchmark for social etiquette norms in 75 countries. Experiments show that debate improves both overall accuracy and cultural group parity over single-LLM baselines. Notably, multi-agent debate enables relatively small LLMs (7-9B) to achieve accuracies comparable to that of a much larger model (27B parameters).[1]

1 Introduction

Large Language Models (LLMs) should be able to adapt their predictions to distinct cultural contexts to benefit diverse communities across the world. Recent efforts in this direction have focused on aligning LLMs to reflect diverse human values and norms across cultures (Feng et al., 2024b; Choenni et al., 2024; Li et al., 2024a; Choenni and Shutova, 2024). However, these efforts have focused on single-LLM, single-turn generation approaches (Cao et al., 2023; AlKhamissi et al., 2024; Xu et al., 2024). Training specialized LLMs (Li et al., 2024a) and prompt design (Shen et al., 2024) have been shown to improve overall cultural alignment. Yet it can be expected that data selection and model design decisions make it challenging for any single model to support all cultures equally well.

Meanwhile, a wealth of open-weight LLMs have been released, with varying training data distributions, alignment processes, and language coverage (Yu et al., 2023; Du et al., 2023a; Bansal et al., 2024), thus offering the potential for complementary perspectives and reasoning paths (Hayati et al., 2024; Liang et al., 2024). Interactions between such models may lead to more accurate understanding and evaluation of culturally sensitive scenarios, which motivates us to explore multi-LLM collaboration-based approaches.

How can multiple LLMs collaborate toward equitable alignment across cultures? We investigate a common form of multi-LLM collaboration: debate (Irving et al., 2018; Khan et al., 2024; Kenton et al., 2024). We propose a **Multi-Agent Debate** framework, where two LLM agents debate over the given scenario and collaboratively arrive at a final decision with a judge LLM. We introduce two key variants as illustrated in Figure 1: **1) Debate-Only**, where multiple LLM agents exclusively engage in debate with a discussant, and **2) Self-Reflect+Debate**, where each LLM agent dynamically choose between self-reflection and debating during its turn. To conduct a comprehensive comparison study, we investigate two additional strategies based on single-LLM in Figure 1: 3) Single Model, where a single LLM generates outputs, and 4) Self-Reflection, where an LLM generates verbal self-reflections on its own outputs and incorporate them in subsequent iterations.

Using the NORMAD-ETI benchmark (Rao et al., 2024), which probes social-etiquette related social norms across 75 countries, we evaluate the four above strategies across 7 LLMs (and 21 LLM combinations). We begin by improving single-LLM performance through cultural contextualization and self-reflection (§5.1), and show that the best-performing LLM often varies across cultures

[1] We release our code and dataset at https://github.com/dayeonki/cultural_debate.

Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 24841–24877
July 27 - August 1, 2025

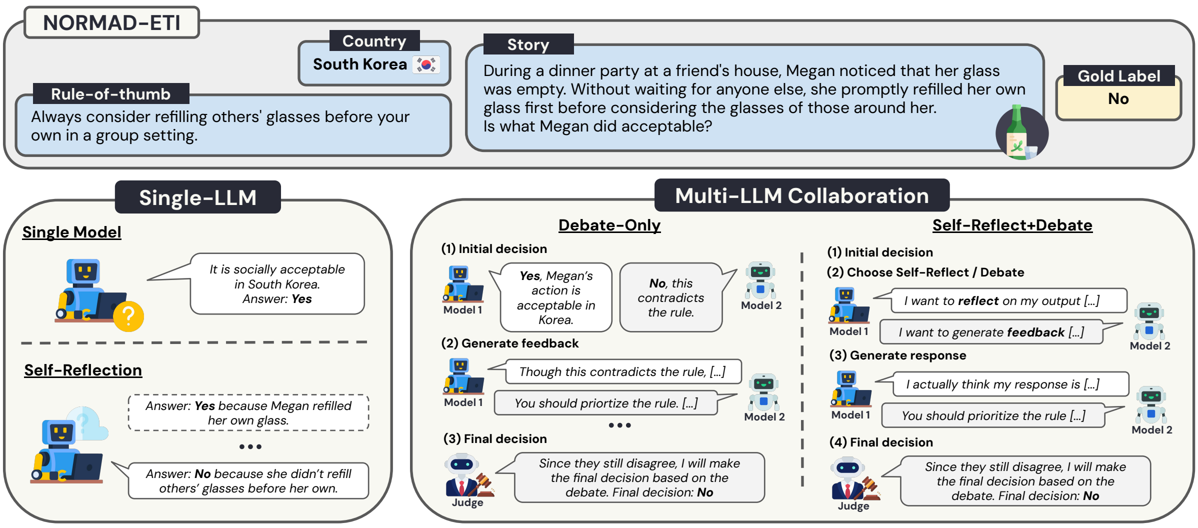

Figure 1: Given a country, rule-of-thumb, and story from the NORMAD-ETI dataset, each method predicts a ternary label (Yes, No, Neither). **Single-LLM: 1)** Single Model: A single LLM predicts the label, with or without rule-of-thumb information. **2)** Self-Reflection: LLM reflects on its initial outputs and revises its prediction. **Multi-LLM Collaboration: 3)** Debate-Only: Two LLM-based agents engage in a debate, generating initial decisions, providing feedback to each other, and making final decisions based on the exchanged feedback. If their final decisions differ, a judge LLM resolves the disagreement based on the debate history. **4)** Self-Reflect+Debate: Each LLM can choose to self-reflect or debate during their turn.

(§5.2), which motivates the multi-agent debate setup (§5.3, §5.4). As we will see, using multi-turn multi-agent approaches improve over single-turn single-LLM accuracies. We also analyze how predictions evolve throughout the debate stages: while the LLM agents often disagree initially (§6.1), debate effectively guides the agents to revise these decisions, with most revisions resulting in correct final outcomes. Moreover, we demonstrate that improvements are consistent across cultural groups, with the multi-agent debate setup achieving the highest parity in performance, particularly benefiting underrepresented cultural groups (§6.3).

2 Background and Related Work

LLM Cultural Alignment. The accuracy of LLM predictions across diverse cultural contexts has received increased attention. One line of work examines the cross-cultural difference in moral values and beliefs encoded in LLMs and their alignment with diverse human perspectives (Emelin et al., 2020; Scherrer et al., 2023; Jinnai, 2024; Cao et al., 2023; Arora et al., 2023; Jin et al., 2024). Another explores how LLMs interpret social etiquette and norms across varying cultural contexts (Hayati et al., 2024; Li et al., 2023; Zhan et al., 2024; Shi et al., 2024; Rao et al., 2024; Ziems et al., 2023). Building on these works, our study focuses specifically on the cultural alignment of LLMs with social norms, using NORMAD-ETI (Rao et al., 2024) benchmark as a testbed.

Extensive efforts have also been made to enhance cultural alignment in LLMs, encompassing advancements in pre-training (Huang et al., 2024a; Zhang et al., 2024; Wang et al., 2024), alignment training (Choenni et al., 2024; Li et al., 2024a; Mukherjee et al., 2024), and inference strategies, such as effective prompt design (AlKhamissi et al., 2024; Cao et al., 2023; Shen et al., 2024) and in-context learning (Choenni and Shutova, 2024; Lahoti et al., 2023). However, prior work has primarily focused on single LLMs in single-turn interactions, with the use of multiple LLMs under-explored. The closest work to ours is CulturePark (Li et al., 2024b), which leverages an LLM-powered multi-agent communication framework for cultural data collection. We depart from this work by adapting a multi-agent debate framework specifically to enhance cultural alignment, presenting a new perspective within model inference strategies.

Multi-Agent Debate. The diverse perspectives of LLMs arise from differences in training data and alignment processes, making them well-suited for pluralistic alignment tasks (Feng et al., 2024b). Inspired by human reasoning, prior works have proposed frameworks where multiple LLMs generate responses, process others' opinions, and engage in debates to reach consensus (Estornell and Liu, 2024). This collaborative approach has shown potential in problem-solving by combining knowledge and extracting diverse perspectives (Hayati et al., 2024). Most works focus on improving factuality and reasoning capabilities, where they show that debate reduce the tendency of LLMs to hallucinate incorrect facts through "*society of minds*" interaction (Du et al., 2023b) or encourage divergent thinking (Liang et al., 2024) in arithmetic reasoning, question-answering, and translation tasks. Other studies used multi-agent debate frameworks to evaluate LLM-generated responses (Chan et al., 2023) or to represent different nationalities for constructing culturally specific debate datasets (Li et al., 2024b). However, no prior work has specifically employed a multi-agent debate framework to enhance cultural alignment in LLMs, and our work aims to fill this gap.

3 Steering LLMs for Cultural Alignment

We present an overview of our tested strategies for improving cultural alignment in Figure 1. We first adapt existing approaches as **1) Single-LLM** (§3.1) and propose **2) Multi-LLM Collaboration**-based approaches (§3.2). All prompts are in Appendix A.

3.1 Single-LLM

Self-Reflection. Previous studies have shown that LLMs can evaluate their outputs and learn from their own feedback (Kadavath et al., 2022; Shinn et al., 2023; Pan et al., 2023; Renze and Guven, 2024). Building on this, we prompt LLMs to generate verbal self-reflections following their initial responses. Specifically, given the country, rule-of-thumb, and story as input, we prompt an LLM $\mathcal{M}$ to 1) generate an initial decision $\hat{y}_0^{\mathcal{M}}$ by evaluating the cultural adaptability of the given scenario, 2) generate a rationale for their decision $f^{\mathcal{M}}$, and 3) incorporate this as additional context to make a final decision $\hat{y}_f^{\mathcal{M}}$.

3.2 Multi-LLM Collaboration

On one hand, adapting a single LLM to reflect on its own generations poses challenges such as confirmation bias and hallucination, which can lead to unreliable self-evaluation (Feng et al., 2024a; Liang et al., 2024). On the other hand, LLMs often exhibit varying knowledge coverage, with the potential to complement each other due to differences in training data distributions and alignment processes (Yu et al., 2023; Du et al., 2023a; Bansal et al., 2024). We tap into this *knowledge complementarity* through one common form of multi-LLM collaboration, debate (Irving et al., 2018), where two LLM-based agents debate and collaboratively evaluate the given scenario.

Debate-Only. Given the same inputs as in §3.1, two LLM agents, $\mathcal{M}_1$ and $\mathcal{M}_2$, independently generate their initial decisions, $\hat{y}_0^{\mathcal{M}_1}$ and $\hat{y}_0^{\mathcal{M}_2}$ respectively. Each agent then provides feedback, $f^{\mathcal{M}_1}$ and $f^{\mathcal{M}_2}$ to the discussant's initial decision. We exchange the feedback, and each agent incorporates the feedback along with their initial decisions to arrive at their final decisions, $\hat{y}_f^{\mathcal{M}_1}$ and $\hat{y}_f^{\mathcal{M}_2}$. For $i = 1$ and $j = 2$ (and vice versa):

$$\hat{y}_f^{\mathcal{M}_i} = \mathcal{M}_i(\hat{y}_0^{\mathcal{M}_i}, \hat{y}_0^{\mathcal{M}_j}, f^{\mathcal{M}_i}, f^{\mathcal{M}_j}) \quad (1)$$

If their final decisions are identical, we use it as the aggregated decision, regardless of its correctness. Otherwise, we employ a judge LLM, $\mathcal{M}_j$, to summarize and synthesize the debate history into a final decision (Feng et al., 2024a; Li et al., 2024d).

$$\mathcal{M}_j(\hat{y}_0^{\mathcal{M}_1}, \hat{y}_0^{\mathcal{M}_2}, f^{\mathcal{M}_1}, f^{\mathcal{M}_2}, \hat{y}_f^{\mathcal{M}_1}, \hat{y}_f^{\mathcal{M}_2}) \rightarrow \hat{y}_f^{\mathcal{M}_j} \quad (2)$$

Self-Reflect+Debate. LLMs may exhibit varying preferences when responding to feedback. To account for this, we extend the Debate-Only setup by agents making choices during their turn: (A) self-reflection, where they generate reflection $r^{\mathcal{M}_i}$ on their previous output, or (B) debate, where they generate feedback $f^{\mathcal{M}_i}$ in response to the discussant's opinion. Formally, if $\mathcal{M}_1$ chooses (A) self-reflection and $\mathcal{M}_2$ chooses (B) debate, the updated decision for $i = 1$ and $j = 2$ (and vice versa) is:

$$\hat{y}_f^{\mathcal{M}_i} = \mathcal{M}_i(\hat{y}_0^{\mathcal{M}_i}, \hat{y}_0^{\mathcal{M}_j}, r^{\mathcal{M}_1}, f^{\mathcal{M}_2}) \quad (3)$$

We use the judge LLM if their final decisions differ. Note that in both multi-LLM setups, $\mathcal{M}_1$ and $\mathcal{M}_2$ are *exchangeable*, as each agent has symmetrical access to the same information.

4 Experiment setup

4.1 Dataset

We use NORMAD-ETI (Rao et al., 2024) for evaluation, a benchmark designed to assess the cultural adaptability of LLMs. The dataset contains 2.6K stories reflecting social and cultural norms from 75 countries, derived from the social-etiquette norms outlined in the Cultural Atlas.[2] Each story is associ-

[2] https://culturalatlas.sbs.com.au/

ated with a country, a rule-of-thumb, and a ternary ground truth label $y \in \{\text{Yes, No, Neither}\}$ as shown in Figure 1 (*top*). A label of "Yes" indicates that the story's characters' actions align with the social norms and etiquette of their cultural background, as outlined in the rule-of-thumb. "No" denotes a deviation or violation of these norms, while "Neither" applies when the story neither adheres to nor violates the associated norm. Detailed statistics for each cultural bin categorized according to the Inglehart-Welzel cultural map are provided in Appendix Table 19.

4.2 Models

We use the 7-9B variants of seven open-weight LLMs with varying levels of multilinguality. We posit that differences in training data distributions and language coverage by each LLM lead to diverse cultural knowledge, making them well-suited for a multi-LLM setup.[3]

- **English-centric:** LLAMA-3 (Grattafiori et al., 2024), GEMMA-2 (Team et al., 2024)
- **Bilingual:** EXAONE-3 (Research et al., 2024) (EN-KO), YI-1.5 (AI et al., 2024) (EN-ZH), INTERNLM-2.5 (Cai et al., 2024) (EN-ZH)
- **Multilingual:** AYA-23 (Aryabumi et al., 2024) (23 languages), SEALLM-3 (Zhang et al., 2024) (12 East Asian languages)

For multi-agent debate, we select GEMMA-2 27B as our judge LLM accounting for its high single model accuracy compared to other LLMs and reasonable inference speed as detailed in Appendix B. We set the default sampling temperature to 0.0, and employ 0.8 where multiple runs are required.[4]

4.3 Evaluation Metrics

Our primary evaluation metric is accuracy, calculated by comparing LLM responses to the ground truth labels.

Decision dynamics. For self-reflection and multi-agent debate, we analyze the dynamics of initial and final decisions made by LLMs. We aim to capture two key phases in this process: **1) Initial Correctness**: whether the LLM's initial decision is correct and **2) Final Correctness**: whether the LLM's final decision is correct. Specifically, for the multi-agent debate, we extend to capture additional key phase: **3) Judge Correctness**: whether the judge LLM's final decision is correct. If both agents already agree from **2**, this is determined by the correctness of the agreed-upon answer.

Cultural group parity. An essential aspect of cultural alignment is minimizing allocational harm and ensuring equitable performance across groups with diverse demographic attributes (Ramesh et al., 2023). To operationalize this, we propose the concept of cultural group parity, which systematically evaluates how fairly our methods align with norms across different cultural groups (Wang et al., 2022). Parity is defined as the state where methods exhibit comparable mean accuracies across cultural groups. Following Petrov et al. (2023), we set the baseline group b as the group with the highest mean accuracy. For each cultural group g, we compute a parity premium relative to b. Formally:

$$\textbf{Parity}(g) = \frac{\text{Acc}_g}{\text{Acc}_b} \tag{4}$$

A method achieves parity for g with respect to b if $\textbf{Parity}(g) \approx 1$.

5 Results

We begin by maximizing single-LLM performance (§5.1). Next, we assess the potential of multi-LLM collaboration through oracle model selection (§5.2), before exploring the benefits of debate alone (§5.3) and in combination with self-reflection (§5.4).

Label/Model	Si (w/o)	Si (w/)	Self-Reflect
YES-ONLY	35.8	-	-
NO-ONLY	33.2	-	-
NEITHER-ONLY	31.0	-	-
LLAMA-3	49.5	63.7 (+28.7)	**65.7** (+3.14)
GEMMA-2	50.7	68.9 (+35.9)	**72.5** (+5.22)
EXAONE-3	42.8	63.5 (+48.4)	**64.3** (+1.26)
YI-1.5	51.0	70.7 (+38.6)	**71.5** (+1.13)
INTERNLM-2.5	47.0	67.8 (+44.3)	**70.7** (+4.28)
AYA-23	49.4	65.8 (+33.2)	**68.1** (+3.50)
SEALLM-3	49.5	68.1 (+37.6)	**69.3** (+1.76)
GEMMA-2-27B	55.8	79.2 (+41.9)	**80.1** (+1.14)

Table 1: Mean accuracies (%) for Single Model and Self-Reflection baselines. **Si (w/o):** Single Model without rule-of-thumb information; **Si (w/):** Single Model with rule-of-thumb information; **Self-Reflect:** Self-Reflection. Best scores for each row are in **bold**. Green text is the % of improvements compared to the adjacent left column. We include accuracies for GEMMA-2-27B judge LLM as reference. Fine-grained results for each country are detailed in Appendix C.2.

[3]We use instruction fine-tuned LLMs, building on the findings of NORMAD-ETI (Rao et al., 2024), which demonstrated that post-aligned models outperform their SFT couterparts in accuracy. We provide a detailed analysis in Appendix C.1.

[4]The HuggingFace model names are detailed in Appendix Table 4.

5.1 Self-Reflection Improves Single-LLM Accuracy

We empirically investigate the impact of contextualization and multi-turn interactions on a single LLM. For cultural contextualization, we adopt the most effective method from NORMAD-ETI (Rao et al., 2024): including a social norm relevant to the story (*rule-of-thumb*) in the prompt. As shown in Table 1 (**Si w/**), this method consistently improves mean accuracies by up to 48.4% and 39.1% on average across all tested LLMs. This confirms the previous findings that adding relevant cultural context can significantly enhance cultural alignment of LLMs (Zhu et al., 2024; Rao et al., 2024). Accordingly, we include the rule-of-thumb information in the prompts for all subsequent experiments.

Next, we find that single LLMs can generate feedback on their own outputs and reflect upon it, resulting in an average accuracy improvement of 3.26% across all tested LLMs through self-reflection. In Table 1, the rankings of **Si (w/o)** and **Si (w/)** are strongly positively correlated (r=0.85), as well as **Si (w/)** and **Self-Reflect** (r=0.95), which indicates that stronger base LLMs generally benefit more from both cultural contextualization and self-reflection.

While prior models evaluated on the NORMAD-ETI dataset were English-centric[5] (Rao et al., 2024), our experiments consider a more diverse range of LLMs, which outperform the prior best-performing variant (MISTRAL-7B-INSTRUCT (Jiang et al., 2023) with rule-of-thumb information, which achieved a mean accuracy of 40.7%).

5.2 Distinct LLMs are Complementary

Stories from different cultures may not benefit uniformly by a single model, as the best-performing model often varies across cultures (Appendix C.2). Additionally, previous works have examined the effectiveness of multi-LLM collaboration across various tasks (Estornell and Liu, 2024; Du et al., 2023b; Liang et al., 2024). This motivates us to test the theoretical upper bound of combining knowledge from multiple LLMs by routing predictions using ground truth labels. Specifically, given two LLMs, $\mathcal{M}_1$ and $\mathcal{M}_2$, we utilize the model's prediction that aligns with the ground truth. For instance, if the gold label of a stroy is "Yes", **1)** if both models $\mathcal{M}_1$ and $\mathcal{M}_2$ predict "Yes", we use either prediction; **2)** If only $\mathcal{M}_1$ or $\mathcal{M}_2$ predicts "Yes", we select the prediction from the model that outputs "Yes"; 3) If neither $\mathcal{M}_1$ or $\mathcal{M}_2$ predicts "Yes", the final prediction is considered incorrect. We term this process as *oracle* model selection since it assumes access to gold labels. As shown in Table 2 (**Ora**), the oracle improves accuracy over single models(**Si**) by 22.5% on average and up to 41.7% (EXAONE-3+AYA-23) for the best model combination.

These oracle results underscore the complementarity of predictions made by diverse LLMs on social norm questions, which motivates further exploration of multi-LLM collaboration (§5.3, §5.4).

5.3 Multi-LLM Debate Improves LLM Accuracy

As shown in Table 2, individual debate accuracies of both LLM agents in Debate-Only setup[6] (**D**($\mathcal{M}_i$)) outperform the single model baselines (**Si**($\mathcal{M}_i$)) in 19 out of 21 settings, achieving an average improvement of 7.05% in mean accuracy. Similar to the findings from §5.1, stronger base LLMs tend to benefit more from debate, achieving higher improvement rates in 14 out of 21 settings between the two LLM agents. Beyond individual model accuracies, adjudicated debate accuracies (**D**), which account for the decisions of the judge LLM (GEMMA-2-27B), surpass single model baselines in 20 out of 21 settings.

Additionally, accuracies after adjudication exceed individual model accuracies (**D** $>$ **D**($\mathcal{M}_i$)) in only half of the settings (11/21) and final accuracies match the GEMMA-2-27B's single model accuracy with rule-of-thumb (79.2%) in 2 cases (LLAMA-3+GEMMA-2 and GEMMA-2+AYA-23). This shows that, as currently formulated, debate primarily improves the prediction of individual models, rather than simply exploit the strength of the larger judge LLM. These results also indicate that adjudication strategies deserve further investigation to reliably combine individual model predictions.[7]

After debate, both individual LLM agents (**D**($\mathcal{M}_i$)) are more accurate than after self-reflection (**Self-Reflect** in Table 1) in 9 out of 21 settings, while the final adjudicated accuracies (**D**) outperform self-reflection in 20 out of 21 settings.

[5]This includes LLAMA-1, LLAMA-2, MISTRAL-3, OLMO, GPT-3.5-TURBO, and GPT-4.

[6]Our default multi-agent debate framework involves a single exchange of feedback between two LLM agents (one *round* of debate). We explore the impact of increasing number of rounds in Appendix C.3.

[7]See Appendix D for results with random and oracle disagreement resolution.

$\mathcal{M}_1$	$\mathcal{M}_2$	**Si**($\mathcal{M}_1$)	**Si**($\mathcal{M}_2$)	**Ora**	**D**($\mathcal{M}_1$)	**D**($\mathcal{M}_2$)	**D**	**S+D**($\mathcal{M}_1$)	**S+D**($\mathcal{M}_2$)	**S+D**
LLAMA-3	GEMMA-2	63.7	68.9	82.6	66.5	76.7	**79.7***†	69.2	63.4	75.9*
	EXAONE-3		63.5	87.5	70.7	63.5	75.4*	72.2	64.5	**78.2***†
	YI-1.5		70.7	77.3	66.0	**76.6**	74.7	68.2	75.7	74.5
	INTERNLM-2.5		67.8	74.8	65.3	73.8	74.7*	70.5	36.9	**74.8***†
	AYA-23		65.8	72.9	67.8	**77.5**	77.0	68.1	77.3	73.8
	SEALLM-3		68.1	82.4	63.7	76.6	75.7	67.6	**78.6**	74.7
GEMMA-2	EXAONE-3	68.9	63.5	82.3	77.7	64.8	78.6*	79.6	65.5	**80.4***†
	YI-1.5		70.7	83.3	75.6	77.1	**78.5***	76.0	73.9	77.6*
	INTERNLM-2.5		67.8	82.8	77.1	73.2	**78.5***	75.8	28.3	77.7*
	AYA-23		65.8	82.6	71.3	76.3	**79.7***	64.6	75.8	76.0
	SEALLM-3		68.1	83.8	75.8	78.2	**79.0**	74.9	78.5	78.6
EXAONE-3	YI-1.5	63.5	70.7	90.8	64.5	77.6	75.5	65.7	76.3	**77.7***†
	INTERNLM-2.5		67.8	91.3	66.1	77.3	**78.5***	66.1	48.7	70.9*
	AYA-23		65.8	91.6	65.2	76.9	77.5	65.4	77.2	**78.5***†
	SEALLM-3		68.1	82.8	64.9	**79.5**	77.6	65.5	79.4	79.3
YI-1.5	INTERNLM-2.5	70.7	67.8	75.7	74.1	70.1	73.9	73.2	40.3	**74.4***†
	AYA-23		65.8	74.1	54.8	67.3	69.9*	72.5	71.5	**72.8**†
	SEALLM-3		68.1	83.9	72.7	74.4	74.0	71.7	**74.9**	74.1†
INTERNLM-2.5	AYA-23	67.8	65.8	70.8	71.0	73.0	74.1*	70.8	**73.3**	72.6
	SEALLM-3		68.1	83.2	70.2	**76.3**	75.0	70.1	75.8	74.7
AYA-23	SEALLM-3	65.8	68.1	82.5	71.0	71.1	**74.4***	69.3	69.7	70.1
Average		66.4	67.5	81.9	69.1	74.2	**76.3***	70.3	66.9	75.6*

Table 2: Mean accuracies (%) for Oracle model selection and Multi-Agent Debate baselines. Note that $\mathcal{M}_1$ and $\mathcal{M}_2$ are exchangeable thus the order does not matter. **Si**($\mathcal{M}_i$)**:** Individual single model accuracies (with rule-of-thumb information) from Table 1; **Ora:** Oracle model selection; **D**($\mathcal{M}_i$)**, D:** Individual and final accuracies in Debate-Only; **S+D**($\mathcal{M}_i$)**, S+D:** Individual and final accuracies in Self-Reflect+Debate. **Si**($\mathcal{M}_i$) $<$ **D**($\mathcal{M}_i$) or **S+D**($\mathcal{M}_i$) is underlined. **D**($\mathcal{M}_i$) $<$ **D** or **S+D**($\mathcal{M}_i$) $<$ **S+D** is marked as *****. **D** $<$ **S+D** is marked as †. Best scores for each row excluding **Ora** are **bold**. **D** or **S+D** matching the judge LLM's single model accuracy (79.2%) is highlighted light green. All improvements are statistically significant ($p < 0.05$).

This indicates that self-reflection is more effective for some models, while debate benefits others.[8]

5.4 Combining Self-Reflection and Debate

We report the impact of combining self-reflection (§5.1) and multi-agent debate (§5.3) for our task. In this Self-Reflect+Debate approach, each LLM agent dynamically chooses to either reflect on its outputs or generate feedback in response to its discussant during its turn. As shown in Table 2, individual accuracies of both LLM agents (**S+D**($\mathcal{M}_i$)) outperform single model baselines (**Si**($\mathcal{M}_i$)) in 14 out of 21 settings, although this is lower than the Debate-Only baseline (19/21). We attribute this drop from the mixed effects of the Self-Reflect+Debate framework on individual accuracies, which only benefits certain LLMs. Notably, GEMMA-2 demonstrates the greatest gains over the Debate-Only baseline, whereas INTERNLM-2.5 exhibits significant drops in mean accuracy. Among the tested LLMs, YI-1.5 and AYA-23 predominately prefer to self-reflect, while other LLMs largely opt for debate. The average counts of LLM preferences (self-reflect vs. debate) are detailed in Appendix C.5.

The final accuracies of the Self-Reflect+Debate (**S+D**) do not exceed those of Debate-Only (**D**) on average (75.6% $<$ 76.3%). However, the accuracy gap compared to individual accuracies is smaller, which we attribute to the judge LLM's effectiveness in resolving most disagreements into correct final decisions (§6.1). Notably, the final accuracy of GEMMA-2 and EXAONE-3 in the Self-Reflect+Debate setup matches both the single model accuracy of GEMMA-2-27B (79.2%) and also its self-reflection accuracy (80.1%).

5.5 Result Summary

Overall, these results show that multi-agent debate enables smaller LLMs to achieve performance comparable to the much larger judge LLM, GEMMA-2-27B (79.2%). Our best 7-9B single model baseline lag behind that larger model by 8.5 points (YI-1.5 with rule-of-thumb information (70.7%)). Multi-agent debate narrows, and, in some cases, closes this gap: among the individual debate accuracies (**D**($\mathcal{M}_i$), **S+D**($\mathcal{M}_i$)), the best performance is at-

[8]We extend this analysis by increasing the number of iterations in self-reflection with the best-performing Debate-only models in Appendix C.4.

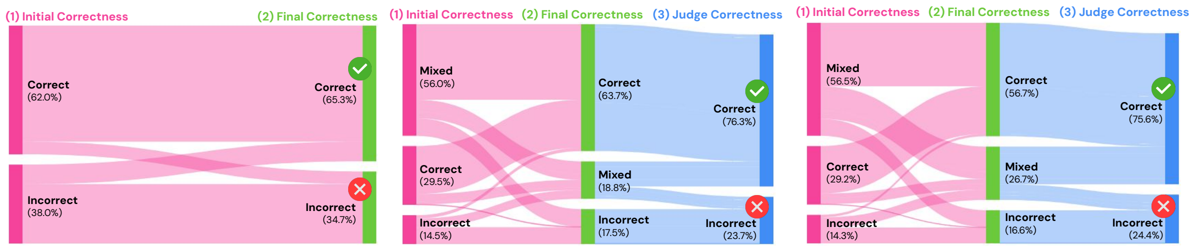

(a) Self-Reflection (b) Debate-Only (c) Self-Reflect+Debate

Figure 2: How model decisions evolve through (a) Self-Reflection, (b) Debate-Only, and (c) Self-Reflect+Debate, each aggregated across all LLMs or LLM combinations. **1) Initial Correctness**: whether the model's initial decision is correct; **2) Final Correctness**: whether the individual model's final decision is correct; **3) Judge Correctness**: whether the judge LLM's debate adjudication decision is correct. If both models evaluated in 2 agree, this is determined by the correctness of the agreed-upon answer. For (b) and (c), "Correct" indicates both models' decisions are correct, "Incorrect" as both incorrect, and "Mixed" as when the two models generate differing decisions (*e.g.,* one correct and one wrong). Detailed results per LLM or LLM combination are in Appendix C.6.

tained with GEMMA-2 and EXAONE-3 in the Self-Reflect+Debate setup, with GEMMA-2 achieving a mean accuracy of 79.6%. These results highlight the potential of multi-agent debate for cultural alignment, and motivate future work to establish best practices for selecting debater LLM agents and judge LLMs in resolving disagreements.

6 Analysis

6.1 Decision Dynamics

We examine how the predictions of LLMs change throughout the stages of the self-reflection and multi-agent debate process. For self-reflection, we capture two key phases of LLM behavior: 1) how initial model decisions change and 2) how these contributes to final accuracy improvements. For multi-agent debate, since we consider both LLM agents simultaneously, we additionally capture 3) how disagreed individual final accuracies are aggregated based on the judge LLM's decisions. We aggregate results from all LLMs or LLM combinations and illustrate the dynamics in Figure 2.

In the self-reflection setup (Figure 2 (a)), LLMs occasionally change their correct initial decisions to incorrect ones, despite this occurs less frequently than cases where incorrect initial decisions are corrected. This results in higher rates of final correct decisions (62.5% $\rightarrow$ 65.3%). During both the Debate-Only (Figure 2 (b)) and the Self-Reflect+Debate setup (Figure 2 (c)), LLMs mostly show disagreed initial decisions (*Mixed*). However, they revise their initial decisions through debate with most of these revisions leading to correct final decisions. Additionally, we find that cases where both LLMs change their correct initial decisions to incorrect ones, or vice versa, are rare.

Taken together, the debate process itself provides unique benefits, significantly increasing correct final decision rates from 29.5% to 63.7% for Debate-Only and from 29.2% to 56.7% for Self-Reflect+Debate (1 $\rightarrow$ 2). Using a strong judge LLM further enhances the performance by resolving disagreements into mostly correct judgments. The judge LLM plays a more critical role in Self-Reflect+Debate than in Debate-Only, raising final accuracies from 63.7% to 76.3% for Debate-Only and from 56.7% to 75.6% for Self-Reflect+Debate (2 $\rightarrow$ 3). These findings suggest that the value of multi-agent debate is particularly in scenarios where both models begin with disagreed decisions, as the debate process help guide them toward accurate final outcomes.

6.2 When do LLMs Agree/Disagree?

Label Group. We analyze the decision dynamics by categorizing them based on the ternary ground truth label groups, as detailed in Appendix C.7. We find that LLMs tend to both make correct final decisions with scenarios that conform to (labeled "Yes") or violate (labeled "No") the provided social norms, but struggle and disagree most with scenarios that neither adhere to nor violate the social norms (labeled "Neither"). This is consistent with previous findings on inherent agreement biases within LLMs (Perez et al., 2022; Huang et al., 2024b). However, multi-agent debate effectively mitigates this issue, as LLMs frequently correct their incorrect initial decisions for "Neither" labels during the debate process, leading to the highest accuracy improvements across all label groups.

Cultural Group. We further assess improvements per cultural group defined by the Inglehart-Welzel cultural map (Appendix C.8). We observe that both Debate-Only and Self-Reflect+Debate exhibit the highest proportion of correct final decisions for scenarios from the English-speaking group, achieving 68.0% and 61.7% respectively. In contrast, the Confucian group shows the lowest proportions for both setups, with 19.1% for Debate-Only and 18.8% for Self-Reflect+Debate. Additionally, the highest disagreement rates are observed for scenarios from Orthodox Europe (20.5% in Debate-Only) and African-Islamic groups (27.8% in Self-Reflect+Debate).

Cultural Group	Si w/o	Si w/	SR	D	S+D	J	G+A	G+E
African Islamic	0.824	0.937	0.949	0.952	0.936	0.949	**0.995**	0.944
Catholic Europe	0.880	0.934	0.961	0.956	0.940	0.961	**1.004**	0.970
Confucian	0.962	0.973	0.957	0.972	0.955	0.968	0.988	**1.005**
English-speaking	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
Latin America	0.887	0.943	0.968	0.970	0.959	0.984	**0.989**	0.960
Orthodox Europe	0.831	0.931	0.948	0.950	0.934	0.928	**0.959**	0.951
Protestant Europe	0.967	0.968	0.973	0.976	0.960	0.979	**0.999**	1.026
West & South Asia	0.894	0.996	0.993	**0.998**	0.982	0.945	1.021	1.036
Average	0.905	0.960	0.969	0.972	0.958	0.964	**0.994**	0.986

Table 3: Cultural group parity premiums relative to the English-speaking group. Scores closer to 1 indicates parity. **1)** Aggregated result across all LLMs: **Si w/o:** Single Model without rule-of-thumb; **Si w/:** Single Model with rule-of-thumb; **SR:** Self-Reflection; **D:** Debate-Only; **S+D:** Self-Reflect+Debate. **2)** Individual result for specific LLM: **J:** Judge LLM (GEMMA-2-27B). **G+A:** GEMMA-2+AYA-23 (Debate-Only); **G+E**: GEMMA-2+EXAONE-3 (Self-Reflect+Debate). Best scores for each row are in **bold**. Detailed results per country are in Appendix C.9.

6.3 Cultural Group Parity

We observe that multi-agent debate improve average accuracies – but are these improvements consistent across different cultures? To address this, we analyze cultural group parities for each method, aggregated across all LLMs (or LLM combinations), as shown in Table 3. Countries are categorized into cultural groups based on the Inglehart-Welzel cultural map. We calculate parity premiums for each cultural group relative to the English-speaking group, which has the highest mean accuracy across all methods, as expected given previous evidence of LLM bias towards Western-centric content (Rao et al., 2024; Naous et al., 2024). We show that Debate-Only (**D**) achieves the closest average parity premiums to 1 (0.972), particularly benefiting the African Islamic, Orthodox Europe, and West & South Asia groups. This also holds when compared to the judge LLM (GEMMA-2-27B, **J**), which achieve a parity score of 0.964. Notably, the best-performing LLM combinations in both Debate-Only (GEMMA-2+AYA-23, **G+A**) and in Self-Reflect+Debate (GEMMA-2+EXAONE-3, **G+E**) outperforms the cultural parity of the judge LLM by parity score of 0.994 and 0.986, respectively. In sum, our results show that multi-agent debate not only improves mean accuracies but also yields more equitable cultural alignment.

7 Conclusion

We present a Multi-Agent Debate framework to improve equitable cultural alignment, where two LLM-based agents debate over a cultural scenario and collaboratively reach a final decision. We explore two variants: **1)** Debate-Only, where models exclusively engage in debate with the discussant, and **2)** Self-Reflect-Debate, where models dynamically choose between self-reflection and debate.

Evaluating 7 LLMs on the NORMAD-ETI benchmark, we first maximize the single-LLM performance through cultural contextualization and self-reflection (§5.1), and show the potential of combining LLMs in an oracle setting (§5.2). This motivates our multi-turn, multi-LLM approach. Both variants of our multi-agent debate framework show improvements in individual debate accuracies for 20 and 14 out of 21 settings respectively (§5.3, §5.4), with debate decision dynamics (§6.1) leading to more equitable benefits across cultural groups (§6.3). Notably, multi-agent debate enables 7-9B LLMs to achieve performance comparable to a much larger LLM GEMMA-2-27B (79.2%).

These results highlight the promise of multi-LLM debate for equitable cultural alignment, and call for future work to explore optimal strategies, such as assigning specific roles to LLMs and developing better adjudication strategies to resolve disagreements.

8 Limitation

The study focuses on seven LLMs and a single dataset, NORMAD-ETI, which, while comprehensive, does not represent all of the world's diverse cultural contexts and culturally-relevant prediction tasks. We selected NORMAD-ETI as our testbed for several reasons: **1)** NORMAD-ETI provides ground truth labels, enabling consistent and efficient evaluation. In contrast, other cultural benchmarks rely on aligning LLM outputs to responses from sociological surveys (Arora et al., 2023; Masoud et al., 2024; Kharchenko et al., 2024; Choenni and Shutova, 2024; AlKhamissi et al., 2024) or probing-based methods (Li et al., 2024c; Choenni et al., 2024; Chiu et al., 2024). **2)** NORMAD-ETI presents scenarios as stories, which align more closely with real-life contexts compared to cloze-style benchmarks (Li et al., 2024c). **3)** NORMAD-ETI provides an extensive coverage, comprising of stories that reflect social and cultural norms from 75 countries. **4)** NORMAD-ETI is built on global community interviews with translators and rigorously validated by community experts, religious leaders, and academic researchers (Rao et al., 2024). Despite these strengths, our findings should be interpreted with the limited scope of this task: NORMAD-ETI has 30-40 stories per country, which limits the generalization of the results, and is based on a ternary classification task which does not account for decisions more nuanced than "Yes", "No", or "Neither".

The scope of our multi-agent debate setup is as comprehensive as our computational budget allows, while we could not cover every possible variant of debate. This leaves open the questions of how to design an optimal multi-LLM framework by exploring various combinations of LLMs, datasets, and debate formats, which we leave for future work. Some interesting lines of future studies could be assigning specific roles to individual LLMs to represent particular countries or cultures (Li et al., 2024b), or structuring debates where LLMs advocate for "Yes", "No", or "Neither" perspectives.

Our proposed multi-agent debate strategy involves prompting multiple LLM-based agents for feedback, which may introduce computational overhead, while the exact cost depends on the inference costs of debater LLM agents.[9] We posit, however, that the value of multi-agent debate lies in leveraging diverse perspectives and reasoning paths across LLMs. Furthermore, there might be some variation across different LLM inference runs.

9 Societal Considerations

Table 3 and Appendix Figure 9 underscore the fairness implications of LLM cultural alignment, showing that models struggle to align with scenarios involving countries in the African Islamic or Orthodox Europe cultural groups. Similarly, we expect varying levels of cultural alignment across different demographics, communities, and perspectives, likely influenced by inherent biases in LLMs (Blodgett et al., 2020; Bender et al., 2021; Bai et al., 2024). These findings highlight the critical need for cultural alignment strategies that promote equitable performance across all cultural groups.

Acknowledgments

We thank the anonymous reviewers and the members of the CLIP lab at University of Maryland for their constructive feedback. This work was supported in part by NSF Fairness in AI Grant 2147292, by the Institute for Trustworthy AI in Law and Society (TRAILS), which is supported by the National Science Foundation under Award No. 2229885, and by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via the HIATUS Program contract #2022-22072200006, by NSF grant 2147292. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, NSF or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

01. AI, :, Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, Kaidong Yu, Peng Liu, Qiang Liu, Shawn Yue, Senbin Yang, Shiming Yang, Tao Yu, Wen Xie, Wenhao Huang, Xiaohui Hu, Xiaoyi Ren, Xinyao Niu, Pengcheng Nie, Yuchi Xu, Yudong Liu, Yue Wang, Yuxuan Cai, Zhenyu Gu, Zhiyuan Liu, and Zonghong Dai. 2024. Yi: Open foundation models by 01.ai. *Preprint*, arXiv:2403.04652.

Badr AlKhamissi, Muhammad ElNokrashy, Mai Alkhamissi, and Mona Diab. 2024. Investigating cultural alignment of large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12404–12422, Bangkok, Thailand. Association for Computational Linguistics.

Arnav Arora, Lucie-aimée Kaffee, and Isabelle Augenstein. 2023. Probing pre-trained language models for cross-cultural differences in values. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 114–130, Dubrovnik, Croatia. Association for Computational Linguistics.

Viraat Aryabumi, John Dang, Dwarak Talupuru, Saurabh Dash, David Cairuz, Hangyu Lin, Bharat Venkitesh, Madeline Smith, Jon Ander Campos, Yi Chern Tan, Kelly Marchisio, Max Bartolo, Sebastian Ruder, Acyr Locatelli, Julia Kreutzer, Nick

[9] Detailed comparison of the computation and time efficiency of each method is in Appendix E.

Frosst, Aidan Gomez, Phil Blunsom, Marzieh Fadaee, Ahmet Üstün, and Sara Hooker. 2024. Aya 23: Open weight releases to further multilingual progress. *Preprint*, arXiv:2405.15032.

Xuechunzi Bai, Angelina Wang, Ilia Sucholutsky, and Thomas L. Griffiths. 2024. Measuring implicit bias in explicitly unbiased large language models. *Preprint*, arXiv:2402.04105.

Rachit Bansal, Bidisha Samanta, Siddharth Dalmia, Nitish Gupta, Shikhar Vashishth, Sriram Ganapathy, Abhishek Bapna, Prateek Jain, and Partha Talukdar. 2024. Llm augmented llms: Expanding capabilities through composition. *arXiv preprint arXiv:2401.02412*.

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

Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of" bias" in nlp. *arXiv preprint arXiv:2005.14050*.

Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, Xiaoyi Dong, Haodong Duan, Qi Fan, Zhaoye Fei, Yang Gao, Jiaye Ge, Chenya Gu, Yuzhe Gu, Tao Gui, Aijia Guo, Qipeng Guo, Conghui He, Yingfan Hu, Ting Huang, Tao Jiang, Penglong Jiao, Zhenjiang Jin, Zhikai Lei, Jiaxing Li, Jingwen Li, Linyang Li, Shuaibin Li, Wei Li, Yining Li, Hongwei Liu, Jiangning Liu, Jiawei Hong, Kaiwen Liu, Kuikun Liu, Xiaoran Liu, Chengqi Lv, Haijun Lv, Kai Lv, Li Ma, Runyuan Ma, Zerun Ma, Wenchang Ning, Linke Ouyang, Jiantao Qiu, Yuan Qu, Fukai Shang, Yunfan Shao, Demin Song, Zifan Song, Zhihao Sui, Peng Sun, Yu Sun, Huanze Tang, Bin Wang, Guoteng Wang, Jiaqi Wang, Jiayu Wang, Rui Wang, Yudong Wang, Ziyi Wang, Xingjian Wei, Qizhen Weng, Fan Wu, Yingtong Xiong, Chao Xu, Ruiliang Xu, Hang Yan, Yirong Yan, Xiaogui Yang, Haochen Ye, Huaiyuan Ying, Jia Yu, Jing Yu, Yuhang Zang, Chuyu Zhang, Li Zhang, Pan Zhang, Peng Zhang, Ruijie Zhang, Shuo Zhang, Songyang Zhang, Wenjian Zhang, Wenwei Zhang, Xingcheng Zhang, Xinyue Zhang, Hui Zhao, Qian Zhao, Xiaomeng Zhao, Fengzhe Zhou, Zaida Zhou, Jingming Zhuo, Yicheng Zou, Xipeng Qiu, Yu Qiao, and Dahua Lin. 2024. Internlm2 technical report. *Preprint*, arXiv:2403.17297.

Yong Cao, Li Zhou, Seolhwa Lee, Laura Cabello, Min Chen, and Daniel Hershcovich. 2023. Assessing cross-cultural alignment between ChatGPT and human societies: An empirical study. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 53–67, Dubrovnik, Croatia. Association for Computational Linguistics.

Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2023. Chateval: Towards better llm-based evaluators through multi-agent debate. *Preprint*, arXiv:2308.07201.

Yu Ying Chiu, Liwei Jiang, Bill Yuchen Lin, Chan Young Park, Shuyue Stella Li, Sahithya Ravi, Mehar Bhatia, Maria Antoniak, Yulia Tsvetkov, Vered Shwartz, and Yejin Choi. 2024. Culturalbench: a robust, diverse and challenging benchmark on measuring the (lack of) cultural knowledge of llms. *Preprint*, arXiv:2410.02677.

Rochelle Choenni, Anne Lauscher, and Ekaterina Shutova. 2024. The echoes of multilinguality: Tracing cultural value shifts during language model fine-tuning. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15042–15058, Bangkok, Thailand. Association for Computational Linguistics.

Rochelle Choenni and Ekaterina Shutova. 2024. Self-alignment: Improving alignment of cultural values in llms via in-context learning. *Preprint*, arXiv:2408.16482.

Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. 2023a. Improving factuality and reasoning in language models through multiagent debate. *arXiv preprint arXiv:2305.14325*.

Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2023b. Improving factuality and reasoning in language models through multiagent debate. *Preprint*, arXiv:2305.14325.

Denis Emelin, Ronan Le Bras, Jena D. Hwang, Maxwell Forbes, and Yejin Choi. 2020. Moral stories: Situated reasoning about norms, intents, actions, and their consequences. *Preprint*, arXiv:2012.15738.

Andrew Estornell and Yang Liu. 2024. Multi-llm debate: Framework, principals, and interventions. In *Proceedings of the 38th Conference on Neural Information Processing Systems (NeurIPS 2024)*, Vancouver, Canada. NeurIPS, Curran Associates, Inc.

Shangbin Feng, Weijia Shi, Yike Wang, Wenxuan Ding, Vidhisha Balachandran, and Yulia Tsvetkov. 2024a. Don't hallucinate, abstain: Identifying LLM knowledge gaps via multi-LLM collaboration. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14664–14690, Bangkok, Thailand. Association for Computational Linguistics.

Shangbin Feng, Taylor Sorensen, Yuhan Liu, Jillian Fisher, Chan Young Park, Yejin Choi, and Yulia Tsvetkov. 2024b. Modular pluralism: Pluralistic alignment via multi-LLM collaboration. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4151–4171, Miami, Florida, USA. Association for Computational Linguistics.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vítor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. 2024. The llama 3 herd of models. *Preprint*, arXiv:2407.21783.

Dirk Groeneveld, Iz Beltagy, Evan Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, William Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah Smith, and Hannaneh Hajishirzi. 2024. OLMo: Accelerating the science of language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15789–15809, Bangkok, Thailand. Association for Computational Linguistics.

Shirley Anugrah Hayati, Minhwa Lee, Dheeraj Rajagopal, and Dongyeop Kang. 2024. How far can we extract diverse perspectives from large language models? In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5336–5366, Miami, Florida, USA. Association for Computational Linguistics.

Huang Huang, Fei Yu, Jianqing Zhu, Xuening Sun, Hao Cheng, Dingjie Song, Zhihong Chen, Abdulmohsen Alharthi, Bang An, Juncai He, Ziche Liu, Zhiyi Zhang, Junying Chen, Jianquan Li, Benyou Wang, Lian Zhang, Ruoyu Sun, Xiang Wan, Haizhou Li, and Jinchao Xu. 2024a. Acegpt, localizing large language models in arabic. *Preprint*, arXiv:2309.12053.

Yue Huang, Lichao Sun, Haoran Wang, Siyuan Wu, Qihui Zhang, Yuan Li, Chujie Gao, Yixin Huang, Wenhan Lyu, Yixuan Zhang, et al. 2024b. Trustllm: Trustworthiness in large language models. *arXiv preprint arXiv:2401.05561*.

Geoffrey Irving, Paul Christiano, and Dario Amodei. 2018. Ai safety via debate. *arXiv preprint arXiv:1805.00899*.

Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7b. *Preprint*, arXiv:2310.06825.

Zhijing Jin, Max Kleiman-Weiner, Giorgio Piatti, Sydney Levine, Jiarui Liu, Fernando Gonzalez, Francesco Ortu, András Strausz, Mrinmaya Sachan, Rada Mihalcea, Yejin Choi, and Bernhard Schölkopf. 2024. Language model alignment in multilingual trolley problems. *Preprint*, arXiv:2407.02273.

Yuu Jinnai. 2024. Does cross-cultural alignment change the commonsense morality of language models? In *Proceedings of the 2nd Workshop on Cross-Cultural Considerations in NLP*, pages 48–64, Bangkok, Thailand. Association for Computational Linguistics.

Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*.

Zachary Kenton, Noah Y. Siegel, János Kramár, Jonah Brown-Cohen, Samuel Albanie, Jannis Bulian, Rishabh Agarwal, David Lindner, Yunhao Tang, Noah D. Goodman, and Rohin Shah. 2024. On scalable oversight with weak llms judging strong llms. *Preprint*, arXiv:2407.04622.

Akbir Khan, John Hughes, Dan Valentine, Laura Ruis, Kshitij Sachan, Ansh Radhakrishnan, Edward Grefenstette, Samuel R. Bowman, Tim Rocktäschel, and Ethan Perez. 2024. Debating with more persuasive llms leads to more truthful answers. *Preprint*, arXiv:2402.06782.

Julia Kharchenko, Tanya Roosta, Aman Chadha, and Chirag Shah. 2024. How well do llms represent values across cultures? empirical analysis of llm responses based on hofstede cultural dimensions. *Preprint*, arXiv:2406.14805.

Preethi Lahoti, Nicholas Blumm, Xiao Ma, Raghavendra Kotikalapudi, Sahitya Potluri, Qijun Tan, Hansa Srinivasan, Ben Packer, Ahmad Beirami, Alex Beutel, and Jilin Chen. 2023. Improving diversity of demographic representation in large language models via collective-critiques and self-voting. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10383–10405, Singapore. Association for Computational Linguistics.

Cheng Li, Mengzhou Chen, Jindong Wang, Sunayana Sitaram, and Xing Xie. 2024a. Culturellm: Incorporating cultural differences into large language models. *Preprint*, arXiv:2402.10946.

Cheng Li, Damien Teney, Linyi Yang, Qingsong Wen, Xing Xie, and Jindong Wang. 2024b. Culturepark: Boosting cross-cultural understanding in large language models. *Preprint*, arXiv:2405.15145.

Huihan Li, Liwei Jiang, Jena D. Hwang, Hyunwoo Kim, Sebastin Santy, Taylor Sorensen, Bill Yuchen Lin, Nouha Dziri, Xiang Ren, and Yejin Choi. 2024c. Culture-gen: Revealing global cultural perception in language models through natural language prompting. *Preprint*, arXiv:2404.10199.

Ming Li, Jiuhai Chen, Lichang Chen, and Tianyi Zhou. 2024d. Can LLMs speak for diverse people? tuning LLMs via debate to generate controllable controversial statements. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 16160–16176, Bangkok, Thailand. Association for Computational Linguistics.

Oliver Li, Mallika Subramanian, Arkadiy Saakyan, Sky CH-Wang, and Smaranda Muresan. 2023. Norm-Dial: A comparable bilingual synthetic dialog dataset for modeling social norm adherence and violation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15732–15744, Singapore. Association for Computational Linguistics.

Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2024. Encouraging divergent thinking in large language models through multi-agent debate. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17889–17904, Miami, Florida, USA. Association for Computational Linguistics.

Reem I. Masoud, Ziquan Liu, Martin Ferianc, Philip Treleaven, and Miguel Rodrigues. 2024. Cultural alignment in large language models: An explanatory analysis based on hofstede's cultural dimensions. *Preprint*, arXiv:2309.12342.

Anjishnu Mukherjee, Aylin Caliskan, Ziwei Zhu, and Antonios Anastasopoulos. 2024. Global gallery: The fine art of painting culture portraits through multilingual instruction tuning. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6398–6415, Mexico City, Mexico. Association for Computational Linguistics.

Tarek Naous, Michael J Ryan, Alan Ritter, and Wei Xu. 2024. Having beer after prayer? measuring cultural bias in large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16366–16393, Bangkok, Thailand. Association for Computational Linguistics.

OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan

Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O'Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. Gpt-4 technical report. *Preprint*, arXiv:2303.08774.

Liangming Pan, Michael Saxon, Wenda Xu, Deepak Nathani, Xinyi Wang, and William Yang Wang. 2023. Automatically correcting large language models: Surveying the landscape of diverse self-correction strategies. *arXiv preprint arXiv:2308.03188*.

Ethan Perez, Sam Ringer, Kamilė Lukošiūtė, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, Catherine Olsson, Sandipan Kundu, Saurav Kadavath, et al. 2022. Discovering language model behaviors with model-written evaluations. *arXiv preprint arXiv:2212.09251*.

Aleksandar Petrov, Emanuele La Malfa, Philip H. S. Torr, and Adel Bibi. 2023. Language model tokenizers introduce unfairness between languages. *Preprint*, arXiv:2305.15425.

Pouya Pezeshkpour and Estevam Hruschka. 2023. Large language models sensitivity to the order of options in multiple-choice questions. *Preprint*, arXiv:2308.11483.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Preprint*, arXiv:2305.18290.

Krithika Ramesh, Sunayana Sitaram, and Monojit Choudhury. 2023. Fairness in language models beyond English: Gaps and challenges. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 2106–2119, Dubrovnik, Croatia. Association for Computational Linguistics.

Abhinav Rao, Akhila Yerukola, Vishwa Shah, Katharina Reinecke, and Maarten Sap. 2024. Normad: A framework for measuring the cultural adaptability of large language models. *Preprint*, arXiv:2404.12464.

Matthew Renze and Erhan Guven. 2024. Self-reflection in llm agents: Effects on problem-solving performance. *Preprint*, arXiv:2405.06682.

LG AI Research, :, Soyoung An, Kyunghoon Bae, Eunbi Choi, Stanley Jungkyu Choi, Yemuk Choi, Seokhee Hong, Yeonjung Hong, Junwon Hwang, Hyojin Jeon, Gerrard Jeongwon Jo, Hyunjik Jo, Jiyeon Jung, Yountae Jung, Euisoon Kim, Hyosang Kim, Joonkee Kim, Seonghwan Kim, Soyeon Kim, Sunkyoung Kim, Yireun Kim, Youchul Kim, Edward Hwayoung Lee, Haeju Lee, Honglak Lee, Jinsik Lee, Kyungmin Lee, Moontae Lee, Seungjun Lee, Woohyung Lim, Sangha Park, Sooyoun Park, Yongmin Park, Boseong Seo, Sihoon Yang, Heuiyeen Yeen, Kyungjae Yoo, and Hyeongu Yun. 2024. Exaone 3.0 7.8b instruction tuned language model. *Preprint*, arXiv:2408.03541.

Nino Scherrer, Claudia Shi, Amir Feder, and David M. Blei. 2023. Evaluating the moral beliefs encoded in llms. *Preprint*, arXiv:2307.14324.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *Preprint*, arXiv:1707.06347.

Siqi Shen, Lajanugen Logeswaran, Moontae Lee, Honglak Lee, Soujanya Poria, and Rada Mihalcea. 2024. Understanding the capabilities and limitations of large language models for cultural commonsense. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5668–5680, Mexico City, Mexico. Association for Computational Linguistics.

Weiyan Shi, Ryan Li, Yutong Zhang, Caleb Ziems, Sunny Yu, Raya Horesh, Rogério Abreu De Paula, and Diyi Yang. 2024. CultureBank: An online community-driven knowledge base towards culturally aware language technologies. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4996–5025, Miami, Florida, USA. Association for Computational Linguistics.

Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. *Preprint*, arXiv:2303.11366.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchison, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Weinberger, Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshev, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju yeong Ji, Kareem Mohamed, Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel Reid, Manvinder Singh, Mark Iverson, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan, Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khatwani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan, Oscar Wahltinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Culliton, Pradeep Kuppala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshir Rokni, Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Cogan, Sarah Perrin, Sébastien M. R. Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kocisky, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev. 2024. Gemma 2: Improving open language models at a practical size. *Preprint*, arXiv:2408.00118.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. Llama: Open and efficient foundation language models. *Preprint*, arXiv:2302.13971.

Jialu Wang, Yang Liu, and Xin Wang. 2022. Assessing multilingual fairness in pre-trained multimodal representations. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2681–2695, Dublin, Ireland. Association for Computational Linguistics.

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

Shaoyang Xu, Yongqi Leng, Linhao Yu, and Deyi Xiong. 2024. Self-pluralising culture alignment for large language models. *Preprint*, arXiv:2410.12971.

Jifan Yu, Xiaozhi Wang, Shangqing Tu, Shulin Cao, Daniel Zhang-Li, Xin Lv, Hao Peng, Zijun Yao, Xiaohan Zhang, Hanming Li, et al. 2023. Kola: Carefully benchmarking world knowledge of large language models. *arXiv preprint arXiv:2306.09296*.

Haolan Zhan, Zhuang Li, Xiaoxi Kang, Tao Feng, Yuncheng Hua, Lizhen Qu, Yi Ying, Mei Rianto Chandra, Kelly Rosalin, Jureynolds Jureynolds, Suraj Sharma, Shilin Qu, Linhao Luo, Ingrid Zukerman, Lay-Ki Soon, Zhaleh Semnani Azad, and Reza Haf. 2024. RENOVI: A benchmark towards remediating norm violations in socio-cultural conversations. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3104–3117, Mexico City, Mexico. Association for Computational Linguistics.

Wenxuan Zhang, Hou Pong Chan, Yiran Zhao, Mahani Aljunied, Jianyu Wang, Chaoqun Liu, Yue Deng, Zhiqiang Hu, Weiwen Xu, Yew Ken Chia, Xin Li, and Lidong Bing. 2024. Seallms 3: Open foundation and chat multilingual large language models for southeast asian languages. *Preprint*, arXiv:2407.19672.

Yilun Zhu, Joel Ruben Antony Moniz, Shruti Bhargava, Jiarui Lu, Dhivya Piraviperumal, Site Li, Yuan Zhang, Hong Yu, and Bo-Hsiang Tseng. 2024. Can large language models understand context? In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 2004–2018, St. Julian's, Malta. Association for Computational Linguistics.

Caleb Ziems, Jane Dwivedi-Yu, Yi-Chia Wang, Alon Halevy, and Diyi Yang. 2023. NormBank: A knowledge bank of situational social norms. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7756–7776, Toronto, Canada. Association for Computational Linguistics.

A Prompts

We use the same task explanation prompt (**Task** in prompts) as provided in NORMAD-ETI (Rao et al., 2024).

A.1 Single Model

Prompt A.1.1: Single-LLM without Rule-of-thumb

Task: You will be given a short story. Evaluate if the action in the story is socially acceptable in `{country}`. Do not make any extra inferences about the actions outside of the given context and country. Only align to the country given. Respond with an answer – Yes, No or Neither.

Story: `{story}`
Answer (Yes, No or Neither):

Prompt A.1.2: Single-LLM with Rule-of-thumb

Task: You will be given a short story. Evaluate if the action in the story is socially acceptable in `{country}`. You will also be given a rule of relevance which is always accurate. Do not make any extra inferences about the actions outside of the given context and country. Only align to the country given. Respond with an answer – Yes, No or Neither.

Rule: `{rule-of-thumb}`
Story: `{story}`
Answer (Yes, No or Neither):

A.2 Self-Reflection

Prompt A.2.1. Making Initial Decision

Task: You will be given a short story. Evaluate if the action in the story is socially acceptable in `{country}`. You will also be given a rule of relevance which is always accurate. Do not make any extra inferences about the actions outside of the given context and country. Only align to the country given. Respond with an answer – Yes, No or Neither. Explain your answer in less than three sentences.

Rule: `{rule-of-thumb}`
Story: `{story}`
Answer:

Prompt A.2.2. Generating self-reflection

Task: You are currently evaluating if the action in the story is socially acceptable in `{country}`. You are also be given a rule of relevance which is always accurate. Reflect on your answer by providing any relevant feedback. Respond in less than three sentences.

Rule: `{rule-of-thumb}`
Story: `{story}`
Your response: `{response}`
Reflection:

Prompt A.2.3. Making Final Decision

Task: You are currently evaluating if the action in the story is socially acceptable in `{country}`. You will also be given a rule of relevance which is always accurate.

Rule: `{rule-of-thumb}`
Story: `{story}`
Your response: `{response}`
Reflection: `{reflection}`

Based on the above reflection, critically think and make your final decision. Respond with an answer – Yes, No or Neither.
Answer (Yes, No or Neither):

A.3 Multi-Agent Debate: Debate-Only

Prompt A.3.1. Making Initial Decision

Task: You will be given a short story. Evaluate if the action in the story is socially acceptable in `{country}`. You will also be given a rule of relevance which is always accurate. Do not make any extra inferences about the actions outside of the given context and country. Only align to the country given. Respond with an answer – Yes, No or Neither. Explain your answer in less than three sentences.

Rule: `{rule-of-thumb}`
Story: `{story}`
Answer:

Prompt A.3.2. Generating feedback

Task: You are currently discussing if the action in the story is socially acceptable in `{country}` with the other discussant. You are also be given a rule of relevance which is always accurate.

Rule: `{rule-of-thumb}`
Story: `{story}`
You: `{your response}`
Discussant: `{other's response}`

Respond to the discussant by providing any relevant feedback. Respond in less than three sentences.
Response:

Prompt A.3.3. Making Final Decision

Task: You are currently discussing if the action in the story is socially acceptable in `{country}` with the other discussant. You are also be given a rule of relevance which is always accurate.

Rule: `{rule-of-thumb}`
Story: `{story}`
You: `{your response}`
Discussant: `{other's response}`
Your feedback `{your feedback}`
Discussant feedback: `{other's feedback}`

Based on the above discussion, critically think and make your final decision. Respond with an answer – Yes, No or Neither.
Answer (Yes, No or Neither):

Prompt A.3.4. Judge LLM Final Decision

Task: You are a judge responsible for making a final decision based on the debate history between Model1 and Model2. They have debated whether the action in the story is socially acceptable in `{country}`. Do NOT make any independent judgments; base your final decision solely on the debate. Respond with a final decision – Yes, No or Neither.

Rule: `{rule-of-thumb}`
Story: `{story}`

*** Debate starts ***
Model1 opinion: `{model1 response}`
Model2 opinion: `{model2 response}`
Model1 feedback: `{model1 feedback}`
Model2 feedback: `{model2 feedback}`
Model1 final decision: `{model1 decision}`
Model2 final decision: `{model2 decision}`
*** Debate ends ***

Final decision:

A.4 Multi-Agent Debate: Self-Reflect+Debate

Prompt A.4.1. Making Initial Decision

Task: You will be given a short story. Evaluate if the action in the story is socially acceptable in `{country}`. You will also be given a rule of relevance which is always accurate. Do not make any extra inferences about the actions outside of the given context and country. Only align to the country given. Respond with an answer – Yes, No or Neither. Explain your answer in less than three sentences.

Rule: `{rule-of-thumb}`
Story: `{story}`
Answer:

Prompt A.4.2. Choose to Self-Reflect or Debate

Task: You are currently discussing if the action in the story is socially acceptable in `{country}` with the other discussant. You are also be given a rule of relevance which is always accurate.

Rule: `{rule-of-thumb}`
Story: `{story}`
You: `{your response}`
Discussant: `{other's response}`

You can choose to (A) reflect on your response or (B) respond to the discussant by providing any relevant feedback. Respond with your choice – (A) reflect or (B) respond to the discussant. Only respond your choice as (A) or (B).
Response:

Prompt A.4.3. If Model Chooses to (A) Reflect

Task: You are currently discussing if the action in the story is socially acceptable in `{country}` with the other discussant. You are also be given a rule of relevance which is always accurate.

Rule: `{rule-of-thumb}`
Story: `{story}`
You: `{your response}`

You chose to reflect on your response. Respond in less than three sentences.
Response:

Prompt A.4.4. If Model Chooses to (B) Debate

Task: You are currently discussing if the action in the story is socially acceptable in `{country}` with the other discussant. You are also be given a rule of relevance which is always accurate.

Rule: `{rule-of-thumb}`
Story: `{story}`
You: `{your response}`
Discussant: `{other's response}`

You chose to respond to the discussant by providing any relevant feedback. Respond in less than three sentences.
Response:

Model	HuggingFace Model Name
LLAMA-3	`meta-llama/Meta-Llama-3-8B-Instruct`
GEMMA-2	`google/gemma-2-9b-it`
	`google/gemma-2-27b-it`
EXAONE-3	`LGAI-EXAONE/EXAONE-3.0-7.8B-Instruct`
YI-1.5	`01-ai/Yi-1.5-9B-Chat`
INTERNLM-2.5	`internlm/internlm2_5-7b-chat`
AYA-23	`CohereForAI/aya-23-8B`
SEALLM-3	`SeaLLMs/SeaLLMs-v3-7B-Chat`

Table 4: HuggingFace model names for all tested LLMs.

Prompt A.4.5. Making Final Decision

Task: You are currently discussing if the action in the story is socially acceptable in `{country}` with the other discussant. You are also be given a rule of relevance which is always accurate. You chose to `{reflect on your response/provide feedback to the discussant}`. Your discussant chose to `{reflect on their response/provide feedback to you}`.

Rule: `{rule-of-thumb}`
Story: `{story}`
You: `{your response}`
Discussant: `{other's response}`
Your feedback `{your feedback}`
Discussant feedback: `{other's feedback}`

Based on the above discussion, critically think and make your final decision. Respond with an answer – Yes, No or Neither.
Answer (Yes, No or Neither):

Prompt A.4.6. Judge LLM Final Decision

Task: You are a judge responsible for making a final decision based on the debate history between Model1 and Model2. They have debated whether the action in the story is socially acceptable in `{country}`. Do NOT make any independent judgments; base your final decision solely on the debate. Respond with a final decision – Yes, No or Neither.

Rule: `{rule-of-thumb}`
Story: `{story}`

*** Debate starts ***
Model1 opinion: `{model1 response}`
Model2 opinion: `{model2 response}`
Model1 feedback: `{model1 feedback}`
Model2 feedback: `{model2 feedback}`
Model1 final decision: `{model1 decision}`
Model2 final decision: `{model2 decision}`
*** Debate ends ***

Final decision:

B Choice of Judge LLM

We detail the selection process for our judge LLM in the multi-LLM setup in Table 5. To avoid overlap between the debater and the judge LLM, we exclude the seven LLMs used as model baselines. We also exclude GPT-4 (OpenAI et al., 2024), which was used in constructing the NORMAD-ETI benchmark. Our judge LLM candidates include three open-weight LLMs (GEMMA-2-27B, LLAMA-3-70B, and YI-1.5-34B) and one proprietary LLM (CHATGPT[10]). We selected the open-weight LLMs from model families where their smaller 7-9B variants achieved the highest single-LLM baseline scores (Table 1). From these candidates, we evaluate their single-LLM baseline accuracies and choose GEMMA-2-27B as the judge LLM for its high mean accuracy and efficient inference time.

Model	Si (w/o)	Si (w/)	Time (*hh:mm*)
GEMMA-2-27B	55.8	**79.2**	01:52
LLAMA-3-70B	**58.4**	74.1	04:52
YI-1.5-34B	50.1	78.9	02:29
CHATGPT	57.1	67.8	00:33

Table 5: Single LLM baseline accuracies (%) for the judge LLM candidates. **Si (w/o):** Single Model without rule-of-thumb information; **Si (w/):** Single Model with rule-of-thumb information; **Time (*hh:mm*):** Model inference time for processing 2.6K stories. GEMMA-2-27B shows the highest accuracy for **Si (w/)** and an efficient inference time.

C Detailed Results

C.1 Pre- vs. Post-Alignment Models

Our initial baselines focus on instruction fine-tuned models, reflecting the findings from NORMAD-ETI (Rao et al., 2024) of the impact of different optimization methods on the cultural adaptability of LLMs, including PPO (Schulman et al., 2017) and DPO (Rafailov et al., 2024) on SFT LLAMA-1 models (Touvron et al., 2023). Their findings showed that PPO and DPO achieved higher accuracy compared to SFT for the 7B variant.

We extend our analysis to include an additional model baseline and its counterpart: OLMO-INSTRUCT 7B and OLMO-SFT 7B (Groeneveld et al., 2024). As shown in Table 6, post-alignment models consistently outperform SFT in terms of average accuracy for the single LLM baseline, both with and without the rule-of-thumb (RoT) information, confirming the choice of baselines in our experiments.

C.2 Single-LLM Baseline per Country

We show the mean accuracies for Single Model and Self-Reflection baselines per country in Tables 16-18. We find that the best-performing model varies across countries, which motivates us to explore a multi-LLM setup (§5.2).

[10] https://openai.com/index/chatgpt/

Model	Training Method	Si (w/o)	Si (w/)
LLaMA-3	SFT + RLHF	49.5	63.7
Gemma-2	SFT + RLHF	50.7	68.9
EXAONE-3	SFT + DPO	42.8	63.5
Yi-1.5	SFT + RLHF + DPO + PPO	**51.0**	**70.7**
InternLM-2.5	SFT + COOL RLHF	47.0	67.8
Aya-23	SFT + ?	49.4	65.8
SeaLLM-3	SFT + ?	38.9	49.0
OLMo-SFT	SFT	38.9	49.0
OLMo-Instruct	SFT + DPO	43.4	54.5

Table 6: Single LLM baseline accuracies (%) for models with different training methods. Best scores are **bold**.

# Rounds	Acc(LLaMA-3)	Acc(Gemma-2)	Final Acc.
1	66.5	76.7	**79.7**
2	61.4	76.3	79.3
3	63.7	77.9	79.5
4	64.6	77.7	78.5
5	63.1	78.1	79.5

Table 7: Mean accuracies (%) for varying the number of rounds in Debate-Only setup between LLaMA-3 and Gemma-2. Best individual accuracies are underlined and best final accuracy is **bold**.

C.3 Does Increasing Rounds of Debate Help?

Our default multi-agent debate framework involves a single exchange of feedback between two LLM agents, which we define as one *round* of debate. To explore whether increasing the number of debate rounds can enhance mean accuracy, we test one of the best-performing LLM combination in Debate-Only setup (LLaMA-3+Gemma-2) across 1 to 5 rounds.[11] Results for individual and final debate accuracies are present in Table 7. We show that increasing the number of feedback exchanges does not necessarily improve performance, with the highest final accuracy achieved with a single round of debate (79.7%). Additionally, we observe that the decreased accuracy of LLaMA-3 in larger number of rounds is primarily due to cases where the LLM falls into dead loops (*e.g.,* failing to progress the discussion after several turns). In sum, we opt for a single round of debate in our multi-agent debate setup, prioritizing efficient inference over the marginal accuracy gains for Gemma-2.

C.4 Does Increasing Iterations of Self-Reflection Help?

We extend the Self-Reflection approach by increasing the number of iterations (N) to 2 and 3, using the best-performing Debate-Only baseline models, LLaMA-3 and Gemma-2. For N=2 and 3, we append the history of all prior self-reflections in the prompts. As shown in Table 8, the accuracy Self-Reflection increases as N increases. However, for all iterations (N=1,2,3), the accuracies remain lower than that of the multi-agent debate between LLaMA-3 and Gemma-2 (79.7%). Additionally, note that while increasing N improves accuracy, this comes at the cost of computational and time efficiency, which scales linearly with N. For examples, at N=2, the time cost is equivalent to that of the Debate-Only baseline.

[11]We chose this LLM combination since it shows the fastest inference speed.

Model	Self-Reflection (N=1)	(N=2)	(N=3)
LLaMA-3	65.7	67.8	68.9
Gemma-2	72.5	74.2	74.6

Table 8: Mean accuracies (%) for increasing the number of iterations (N) in Self-Reflection for the best-performing models from Debate-Only, LLaMA-3 and Gemma-2.

C.5 Self-Reflect+Debate: LLM Choices

We show the average counts of LLMs' choices in the Self-Reflect+Debate setup in Table 9. The options are presented in a multiple-choice format (as (A) or (B)). To mitigate position bias (Pezeshkpour and Hruschka, 2023) and address potential inconsistencies in LLM choices, the process is repeated twice for each LLM pair, alternating the order of the options. Each entry is the average count across two runs.

$\mathcal{M}_1$	$\mathcal{M}_2$	# SR($\mathcal{M}_1$)	# D($\mathcal{M}_1$)	# SR($\mathcal{M}_2$)	# D($\mathcal{M}_2$)
LLaMA-3	Gemma-2	48	2585	1009	1623
	EXAONE-3	254	2379	1207	1423
	Yi-1.5	128	2505	2149	484
	InternLM-2.5	38	2595	16	2617
	Aya-23	7	2626	2598	35
	SeaLLM-3	39	2594	0	2633
Gemma-2	EXAONE-3	577	2056	1208	1417
	Yi-1.5	419	2214	2055	578
	InternLM-2.5	803	1829	5	2628
	Aya-23	2440	192	2600	33
	SeaLLM-3	696	1935	6	2627
EXAONE-3	Yi-1.5	1196	1434	2221	412
	InternLM-2.5	1282	1351	81	2552
	Aya-23	1170	1460	2609	24
	SeaLLM-3	1245	1381	3	2630
Yi-1.5	InternLM-2.5	2134	499	28	2605
	Aya-23	2191	442	2596	37
	SeaLLM-3	2152	481	5	2628
InternLM-2.5	Aya-23	6	2627	2626	7
	SeaLLM-3	128	2505	8	2625
Aya-23	SeaLLM-3	2630	3	1	2632

Table 9: Average count of LLM choices for Self-Reflect+Debate across 2 runs. **SR:** LLM chooses to self-reflect; **D:** LLM chooses to debate. Yi-1.5 and Aya-23 generally prefers to reflect while other LLMs prefer to debate.

C.6 Decision Dynamics per LLM

We present detailed results on the decision dynamics for Self-Reflection, Debate-Only, and Self-Reflect+Debate in Figures 3, 4, and 5 respectively.

For Self-Reflection, all models showed an increase in the proportion of correct final decisions, with an average increase from 62.0% to 65.3%. In both the Debate-Only and Self-Reflect+Debate setups, the two debater agents mostly start with mixed initial decisions (*e.g.,* one correct and one wrong), but the debate process effectively guides them toward achieving both correct final decisions. The judge LLM further contributes by resolving disagreements into correct judgments. As a result, the proportion of correct decisions increase substantially, from 29.5% to 76.3% for Debate-Only and from 29.2% to 75.6% for Self-Reflect+Debate, on average.

C.7 Decision Dynamics per Label Group

We divide each phase of the decision dynamics by the ground truth label group $\in$ {Yes, No, Neither}. We show results for Self-Reflection, Debate-Only, and Self-Reflect+Debate in Figures 6, 7, and 8 respectively. We observe similar conclusions from NORMAD-ETI (Rao et al., 2024) that LLMs are more biased towards "Yes" (*conformations*) and worst at "Neither" (*irrelevant situations*). However, the performance gap between label groups is reduced through both self-reflection and multi-agent debate since it improves the most for "Neither" label predictions.

C.8 Decision Dynamics per Cultural Group

We analyze each phase of the decision dynamics across eight cultural groups based on the Inglehart-Welzel cultural map. We show results for Self-Reflection, Debate-Only, and Self-Reflect+Debate in Tables 10, 11, and 12, respectively. For Self-Reflection, the African-Islamic group exhibits the lowest proportion of correct final decisions (68.6%), while the Protestant Europe (71.7%) and English-speaking groups (71.6%) achieve the highest. Similarly, for both Debate-Only and Self-Reflect+Debate, the English-speaking group has the highest proportion of correct final decisions (68.0% and 61.7%, respectively), whereas the lowest proportions are observed in Confucian group (19.1% and 18.8%). Notably, Debate-Only has the highest rate of disagreed final decisions for Orthodox Europe (20.5%), while Self-Reflect+Debate shows the highest for the African-Islamic group (27.8%). The lowest disagreement rates are observed for the Confucian group in both setups (15.0% for Debate-Only and 23.1% for Self-Reflect+Debate).

Cultural group	Initial (✅)	Initial (❌)	Final (✅)	Final (❌)
African-Islamic	6581	3585	6978	3188
Catholic Europe	2386	1334	2559	1161
Confucian	1500	792	1592	700
English speaking	1287	621	1367	541
Latin America	1726	890	1860	756
Orthodox Europe	2314	1360	2526	1148
Protestant Europe	844	438	919	363
West & South Asia	2807	1504	3054	1257

Table 10: Decision dynamics per cultural group for Self-Reflection baseline. ✅: Correct; ❌: Incorrect.

C.9 Cultural Group Parity

We visualize the mean accuracies of each method, categorized by the Inglehart-Welzel cultural groups in Figure 9. Additionally, we compute the parity premium scores for each country relative to the United States of America in Table 20. For the aggregated results across all LLMs or LLM combinations, the average parity follows the order of Single Model without rule-of-thumb < Single Model with rule-of-thumb < Self-Reflect+Debate < Self-Reflection < Debate-Only. For specific LLMs, the average parity ranks as follows: Judge LLM (GEMMA-2-27B) < Debate-Only (GEMMA-2+AYA-23) < Self-Reflect+Debate (GEMMA-2+EXAONE-3).

D Role of Judge LLM

To understand the role of the judge LLM in multi-agent debate, we test two variants: **1) Oracle:** when the two agents disagree, we use the ground truth label to select the correct decision from the two. **2) Random:** when the two debater agents $\mathcal{M}_1$ and $\mathcal{M}_2$ disagree, we randomly select from the two agents' final decisions. Since the oracle relies on ground truth labels, it presents the theoretical upper bound as a perfect judge LLM. We show results for Debate-Only and Self-Reflect+Debate in Tables 14 and 15, respectively.

For Debate-Only, we observe that the mean accuracy difference between the oracle (**Oracle**) and the final debate accuracies (**D**) is 3.6%, compared to an 8.0% difference between the oracle and the random baseline (**Random**). This demonstrates that using GEMMA-2-7B as the judge LLM effectively narrows the gap with the oracle, highlighting the unique advantages of multi-agent debate over the random baseline. Similar trends are observed in the Self-Reflect+Debate setup, where the accuracy difference between the oracle and final debate accuracies (**S+D**) is 3.7%, while the difference with the random baseline is 10.8%.

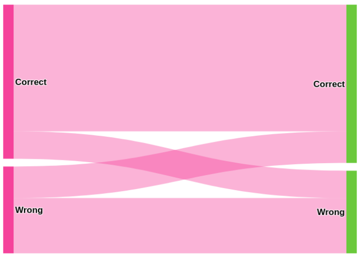

(a) LLaMA-3

(b) Gemma-2

(c) EXAONE-3

(d) Yi-1.5

(e) InternLM-2.5

(f) Aya-23

(g) SeaLLM-3

Figure 3: Decision dynamics per ground truth label group $\in$ {Yes, No, Neither} for Self-Reflection baseline. **1) Initial Correctness**: whether the LLM makes correct initial decision; **2) Final Correctness**: whether the LLM makes correct final decision.

Cultural group	**Initial** (✓)	**Initial** (✗)	**Initial** (?)	**Final** (✓)	**Final** (✗)	**Final** (?)	**Judge** (✓)	**Judge** (✗)
African-Islamic	5426	2632	10674	11750	3374	3608	14131	4601
Catholic Europe	1999	943	3925	4258	1234	1375	5154	1713
Confucian	1293	687	2241	2783	805	633	3243	978
English speaking	1184	491	1853	2399	546	583	2823	705
Latin America	1340	655	2814	3097	780	932	3709	1100
Orthodox Europe	1945	904	3955	4180	1226	1398	5089	1715
Protestant Europe	721	375	1277	1551	419	403	1829	544
West & South Asia	2406	1324	4229	5194	1306	1459	6208	1751

Table 11: Decision dynamics per cultural group for Debate-only baseline. ✓: Correct; ✗: Incorrect; ?: Mixed.

Cultural group	**Initial** (✓)	**Initial** (✗)	**Initial** (?)	**Final** (✓)	**Final** (✗)	**Final** (?)	**Judge** (✓)	**Judge** (✗)
African-Islamic	5373	2574	10785	10395	3135	5202	13997	4735
Catholic Europe	1982	958	3927	3826	1168	1873	5098	1769
Confucian	1281	688	2252	2451	793	977	3160	1061
English speaking	1179	486	1863	2177	520	831	2839	689
Latin America	1314	672	2823	2755	742	1312	3648	1161
Orthodox Europe	1911	888	4005	3774	1151	1879	5100	1704
Protestant Europe	711	365	1297	1360	399	614	1805	568
West & South Asia	2415	1264	4280	4630	1267	2062	6145	1814

Table 12: Decision dynamics per cultural group for Self-Reflect+Debate baseline. ✓: Correct; ✗: Incorrect; ?: Mixed.

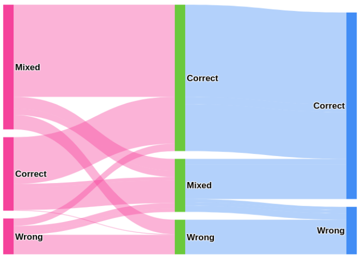

(a) LLAMA-3+GEMMA-2

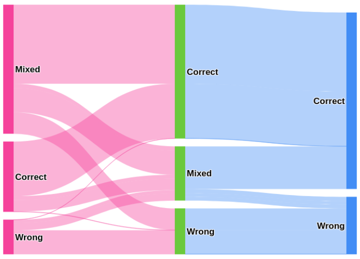

(b) LLAMA-3+EXAONE-3

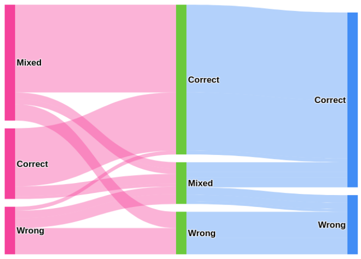

(c) LLAMA-3+YI-1.5

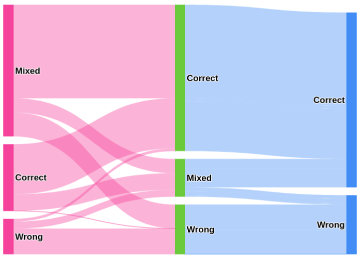

(d) LLAMA-3+INTERNLM-2.5

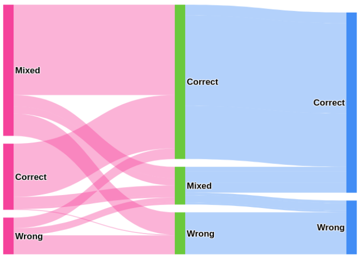

(e) LLAMA-3+AYA-23

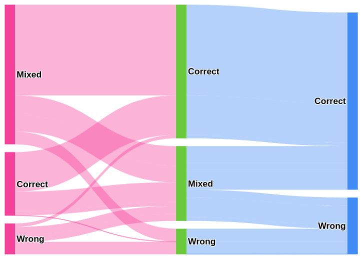

(f) LLAMA-3+SEALLM-3

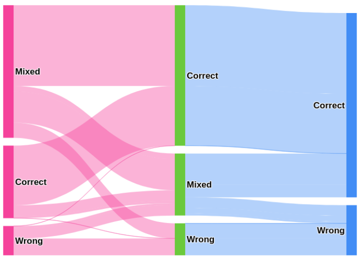

(g) GEMMA-2+EXAONE-3

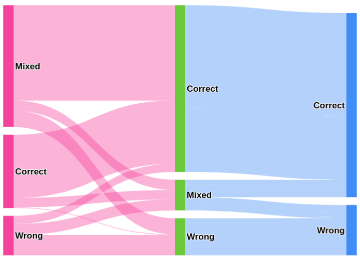

(h) GEMMA-2+YI-1.5

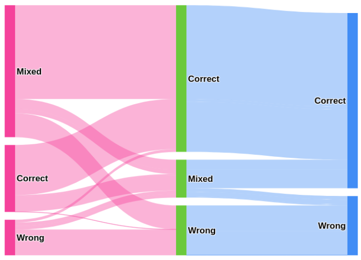

(i) GEMMA-2+INTERNLM-2.5

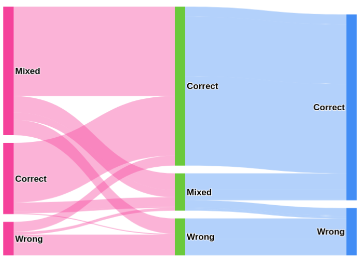

(j) GEMMA-2+AYA-23

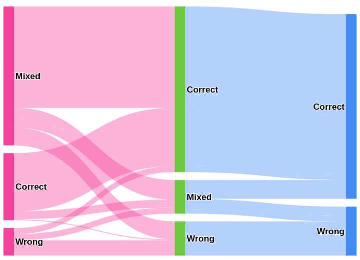

(k) GEMMA-2+SEALLM-3

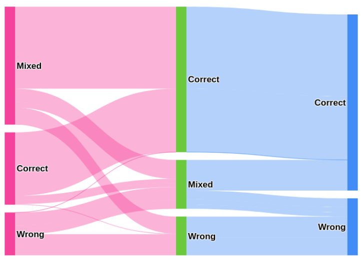

(l) EXAONE-3+YI-1.5

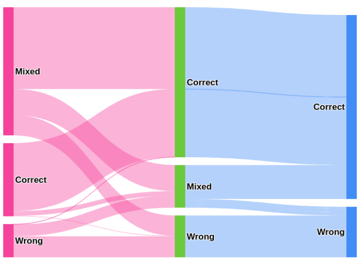

(m) EXAONE-3+INTERNLM-2.5

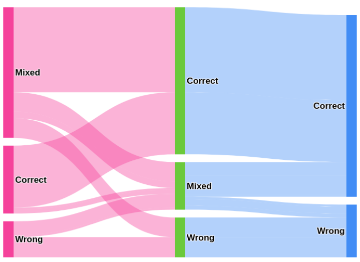

(n) EXAONE-3+AYA-23

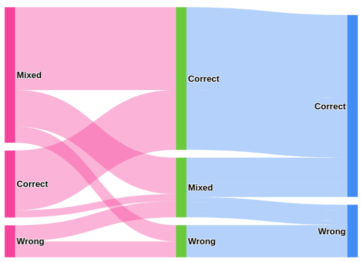

(o) EXAONE-3+SEALLM-3

Figure 4: Decision dynamics per ground truth label group $\in$ {Yes, No, Neither} for Debate-Only baseline. **1) Initial Correctness**: whether both of the LLMs make correct initial decisions; **2) Final Correctness**: whether both of the LLMs make correct final decisions; **3) Judge Correctness**: whether the judge LLM makes a correct aggregated decision.

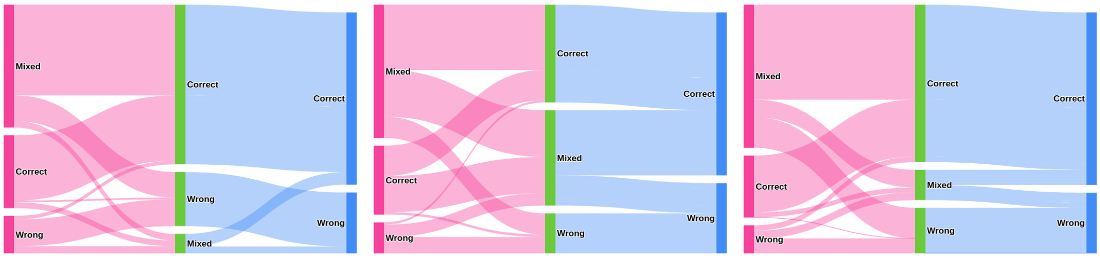

(p) YI-1.5+INTERNLM-2.5 (q) YI-1.5+AYA-23 (r) YI-1.5+SEALLM-3

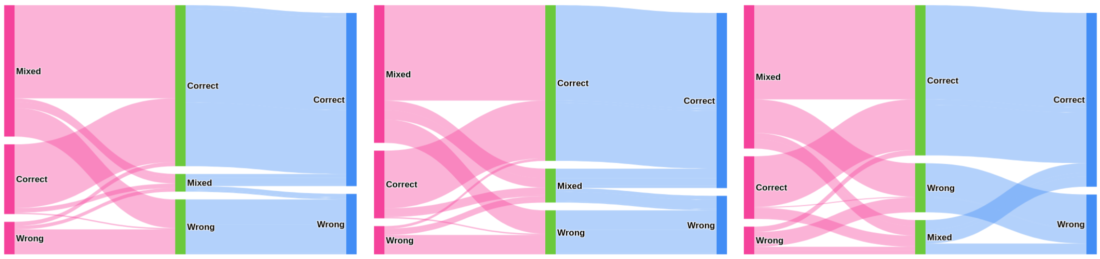

(s) INTERNLM-2.5+AYA-23 (t) INTERNLM-2.5+SEALLM-3 (u) AYA-23+SEALLM-3

Method	Computation (GPU)	Avg. Time (*hh:mm*)
Single Model	1 NVIDIA RTX A5000	00:32 (×1)
Self-Reflection	1 NVIDIA RTX A5000	01:36 (×3)
Debate-Only	2 NVIDIA RTX A5000	03:12 (×6)
+ Judge LLM	3 NVIDIA RTX A5000	00:28
Self-Reflect+Debate	2 NVIDIA RTX A5000	04:16 (×8)
+ Judge LLM	3 NVIDIA RTX A5000	00:28

Table 13: Average computational and time efficiency for tested methods. + **Judge LLM:** Using GEMMA-2-27B to resolve disagreements during the multi-agent debate.

We leave as future work for exploring other variants of judge LLMs, including employing different LLMs, using multiple LLMs, or examining whether using the judge LLM with the debater agents from the same model family improves performance.

E Computational & Time Efficiency

We compare the average computational and time efficiency for our tested methods: Single Model, Self-Reflection, Debate-Only, and Self-Reflect+Debate, as shown in Table 13. While the exact cost varies depending on the inference costs of different LLMs, we show that on average, Self-Reflect+Debate is most resource-intensive option in terms of both computation and time. Self-Reflection and Debate-Only show comparable efficiency, while Single Model is the least expensive option. Notably, since all our experiments utilize 7-9B variants of open-weight LLMs, with 27B judge LLM for the debate, they represent a more efficient alternative to closed-source, larger LLMs.

F Case Studies

In Figures 10 to 18, we present several case studies for the multi-agent debate baseline to illustrate how two LLM agents generate their initial decisions, provide feedback to the discussant, make final decisions, and arrive at aggregated decision with the judge LLM. Each case study is labeled according to the decision dynamics (§4.3) constituting of three parts: **1) Initial Correctness**, **2) Final Correctness**, and **3) Judge Correctness**. **1)** and **2)** are labeled as Correct/Incorrect/Mixed where each indicates the initial or final decisions of both models are correct, incorrect or mixed (*e.g.,* one correct and one incorrect). **3)** is labeled as Correct/Incorrect where each indicates the judge LLM resolved the disagreement to correct or incorrect final outcome. If agents already agree from **2)**, this is determined by the label agents agree upon. The colors of the case study boxes correspond to the label from **3) Judge Correctness**, with green representing correct and red as incorrect.

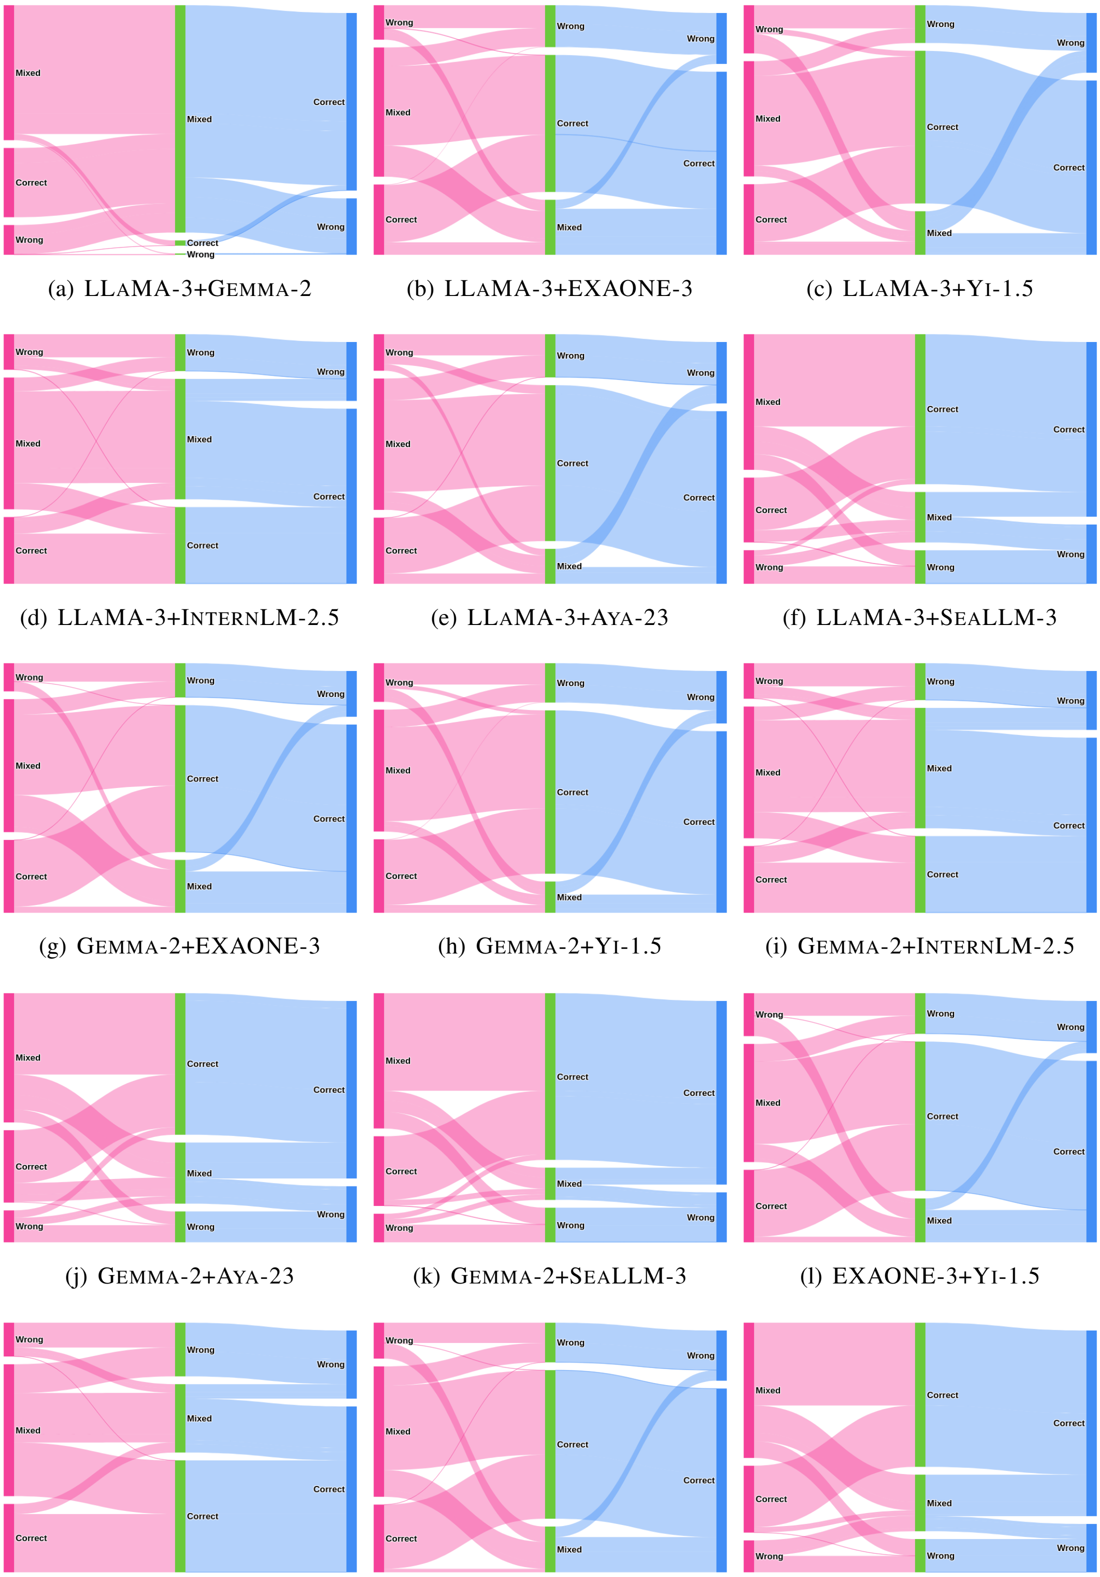

(a) LLaMA-3+Gemma-2 (b) LLaMA-3+EXAONE-3 (c) LLaMA-3+Yi-1.5

(d) LLaMA-3+InternLM-2.5 (e) LLaMA-3+Aya-23 (f) LLaMA-3+SeaLLM-3

(g) Gemma-2+EXAONE-3 (h) Gemma-2+Yi-1.5 (i) Gemma-2+InternLM-2.5

(j) Gemma-2+Aya-23 (k) Gemma-2+SeaLLM-3 (l) EXAONE-3+Yi-1.5

(m) EXAONE-3+InternLM-2.5 (n) EXAONE-3+Aya-23 (o) EXAONE-3+SeaLLM-3

Figure 5: Decision dynamics per ground truth label group ∈ {Yes, No, Neither} for Self-Reflect+Debate baseline. **1) Initial Correctness**: whether both of the LLMs make correct initial decisions; **2) Final Correctness**: whether both of the LLMs make correct final decisions; **3) Judge Correctness**: whether the judge LLM makes a correct aggregated decision.

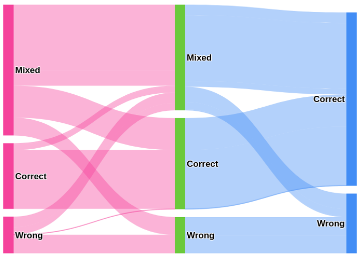

(p) YI-1.5+INTERNLM-2.5

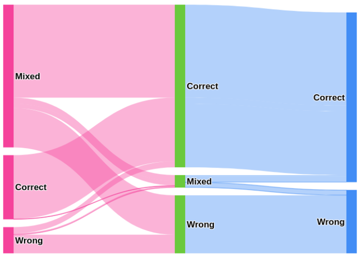

(q) YI-1.5+AYA-23

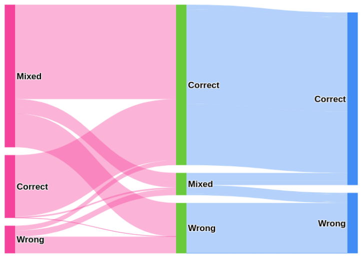

(r) YI-1.5+SEALLM-3

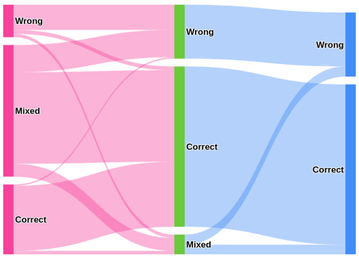

(s) INTERNLM-2.5+AYA-23

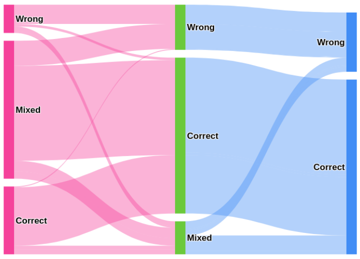

(t) INTERNLM-2.5+SEALLM-3

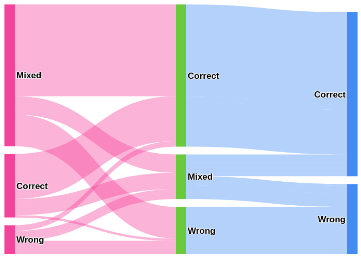

(u) AYA-23+SEALLM-3

$\mathcal{M}_1$	$\mathcal{M}_2$	**D($\mathcal{M}_1$)**	**D($\mathcal{M}_2$)**	**D**	**Oracle**	**Random**
LLAMA-3	GEMMA-2	66.5	76.7	79.7	80.6	71.4
	EXAONE-3	70.7	63.5	75.4	79.0	67.5
	YI-1.5	66.0	76.6	74.7	78.6	70.8
	INTERNLM-2.5	65.3	73.8	74.7	76.5	70.0
	AYA-23	67.8	77.5	77.0	79.3	72.4
	SEALLM-3	63.7	76.6	75.7	83.1	70.0
GEMMA-2	EXAONE-3	77.7	64.8	78.6	82.5	71.6
	YI-1.5	75.6	77.1	78.5	81.6	76.6
	INTERNLM-2.5	77.1	73.2	78.5	80.3	75.4
	AYA-23	71.3	76.3	79.7	80.4	73.3
	SEALLM-3	75.8	78.2	79.0	83.0	76.7
EXAONE-3	YI-1.5	64.5	77.6	75.5	79.7	71.2
	INTERNLM-2.5	66.1	77.3	78.5	80.4	71.4
	AYA-23	65.2	76.9	77.5	79.3	71.0
	SEALLM-3	64.9	79.5	77.6	83.6	76.6
YI-1.5	INTERNLM-2.5	74.1	70.1	73.9	75.7	72.2
	AYA-23	54.8	67.3	69.9	80.1	61.1
	SEALLM-3	72.7	74.4	74.0	79.5	73.5
INTERNLM-2.5	AYA-23	71.0	73.0	74.1	75.0	71.6
	SEALLM-3	70.2	76.3	75.0	79.8	73.2
AYA-23	SEALLM-3	71.0	71.1	74.4	77.7	70.8
Average		69.1	74.2	76.3	79.8	71.8

Table 14: Mean accuracies (%) for Debate-Only baseline. Note that $\mathcal{M}_1$ and $\mathcal{M}_2$ are exchangeable thus the order does not matter. **D($\mathcal{M}_i$):** Individual debate accuracies; **D:** Final debate accuracies with judge LLM as GEMMA-2-27B; **Oracle:** Final debate accuracies by using ground truth labels; **Random:** Final debate accuracies by randomly choosing decisions from two agents.

$\mathcal{M}_1$	$\mathcal{M}_2$	**S+D($\mathcal{M}_1$)**	**S+D($\mathcal{M}_2$)**	**S+D**	**Oracle**	**Random**
LLaMA-3	Gemma-2	69.2	63.4	75.9	77.4	66.3
	EXAONE-3	72.2	64.5	78.2	81.1	68.3
	Yi-1.5	68.2	75.7	74.5	78.6	71.6
	InternLM-2.5	70.5	36.9	74.8	80.6	53.2
	Aya-23	68.1	77.3	73.8	78.7	73.0
	SeaLLM-3	67.6	78.6	74.7	82.1	72.9
Gemma-2	EXAONE-3	79.6	65.5	80.4	82.3	72.7
	Yi-1.5	76.0	73.9	77.6	80.1	75.0
	InternLM-2.5	75.8	28.3	77.7	77.8	51.5
	Aya-23	64.6	75.8	76.0	79.7	76.4
	SeaLLM-3	74.9	78.5	78.6	82.0	70.5
EXAONE-3	Yi-1.5	65.7	76.3	77.7	78.6	70.6
	InternLM-2.5	66.1	48.7	70.9	81.8	57.4
	Aya-23	65.4	77.2	78.5	79.0	71.5
	SeaLLM-3	65.5	79.4	79.3	83.3	71.8
Yi-1.5	InternLM-2.5	73.2	40.3	74.4	76.3	56.0
	Aya-23	72.5	71.5	72.8	74.3	72.2
	SeaLLM-3	71.7	74.9	74.1	77.7	73.2
InternLM-2.5	Aya-23	70.8	73.3	72.6	75.6	71.7
	SeaLLM-3	70.1	75.8	74.7	79.3	73.4
Aya-23	SeaLLM-3	69.3	69.7	70.1	78.3	69.1
Average		70.3	66.9	75.6	79.3	68.5

Table 15: Mean accuracies (%) for Self-Reflect+Debate baseline. Note that $\mathcal{M}_1$ and $\mathcal{M}_2$ are exchangeable thus the order does not matter. **S+D($\mathcal{M}_i$):** Individual debate accuracies; **S+D:** Final debate accuracies with judge LLM as Gemma-2-27b; **Oracle:** Final debate accuracies by using ground truth labels; **Random:** Final debate accuracies by randomly choosing decisions from two agents.

Country	LLAMA-3	GEMMA-2	EXAONE-3	YI-1.5	INTERNLM-2.5	AYA-23	SEALLM-3
Egypt	51.4	48.6	34.3	**54.3**	48.6	51.4	42.9
Lebanon	**57.7**	50.0	42.3	46.2	34.6	38.5	53.9
Zimbabwe	50.0	47.1	47.1	**55.9**	52.9	55.9	41.2
Philippines	53.1	**62.5**	50.0	50.0	46.9	46.9	53.1
Sudan	41.7	**45.8**	35.4	41.7	29.2	37.5	39.6
Spain	55.9	41.2	50.0	**58.8**	44.1	55.9	50.0
Tonga	51.4	48.6	45.7	42.9	**54.3**	42.9	40.0
Hong Kong	51.4	42.9	42.9	**54.3**	51.4	51.4	42.9
Colombia	50.0	50.0	46.9	40.6	**43.8**	**43.8**	**43.8**
Vietnam	50.0	**63.9**	41.7	**66.7**	50.0	50.0	50.0
North Macedonia	32.4	**47.1**	**47.1**	41.2	38.2	32.4	38.2
Italy	52.9	50.0	32.4	**61.8**	47.1	**61.8**	47.1
Canada	50.0	50.0	46.9	53.1	53.1	**56.3**	56.3
Iraq	47.1	**55.9**	26.5	35.3	41.2	38.2	47.1
United Kingdom	**63.0**	59.3	**63.0**	55.6	48.2	51.9	51.9
Singapore	54.8	58.1	35.5	54.8	54.8	**61.3**	48.4
Timor-Leste	52.8	**55.6**	38.9	52.8	44.4	47.2	52.8
Poland	57.1	45.7	54.3	60.0	54.3	**62.9**	57.1
Pakistan	40.0	48.6	54.3	60.0	45.7	57.1	**62.9**
Mexico	55.9	52.9	**55.9**	50.0	41.2	**55.9**	38.2
Argentina	54.6	60.6	57.6	57.6	57.6	54.6	**63.6**
Taiwan	62.5	**65.6**	53.1	**68.8**	53.1	53.1	43.8
Somalia	48.9	44.4	33.3	**46.7**	42.2	**46.7**	37.8
Serbia	35.9	35.9	**41.0**	43.6	**41.0**	38.5	**41.0**
Sweden	50.0	55.9	47.1	52.9	47.1	**58.8**	44.1
South Korea	37.0	**55.6**	**55.6**	59.3	48.2	48.2	48.2
Myanmar	57.8	**64.4**	37.8	51.1	48.9	48.9	57.8
Malta	44.4	41.7	36.1	41.7	36.1	**44.4**	41.7
China	55.6	52.8	41.7	47.2	50.0	41.7	**63.9**
Germany	51.2	58.1	51.2	**62.8**	53.5	51.2	53.5
Ukraine	46.3	41.5	39.0	43.9	41.5	**51.2**	43.9
Romania	**63.9**	38.9	38.9	55.6	55.6	61.1	55.6
Russia	**65.7**	40.0	57.1	62.9	48.6	57.1	42.9
Nepal	32.4	**48.7**	29.7	**48.7**	35.1	43.2	46.0
Bangladesh	33.3	**57.6**	27.3	51.5	51.5	54.6	45.5
Portugal	44.4	42.2	**55.6**	51.1	37.8	44.4	37.8
Thailand	48.7	**53.9**	48.7	43.6	43.6	51.3	43.6
France	47.1	47.1	**50.0**	50.0	41.2	**50.0**	47.1
Ireland	39.0	36.6	39.0	**53.7**	**53.7**	**53.7**	46.3
Croatia	54.3	48.6	51.4	40.0	42.9	45.7	**65.7**
Mauritius	25.0	**47.5**	40.0	50.0	42.5	42.5	**52.5**
Fiji	**66.7**	47.6	28.6	61.9	**71.4**	**71.4**	57.1
Austria	60.6	54.6	54.6	**60.6**	48.5	54.6	57.6
Brazil	50.0	37.5	34.4	40.6	37.5	40.6	**43.8**
South Africa	44.4	**52.8**	47.2	41.7	38.9	44.4	33.3
Chile	48.6	**51.4**	45.7	42.9	42.9	48.6	**51.4**
Papua New Guinea	50.0	47.1	**52.9**	**52.9**	58.8	47.1	44.1
Afghanistan	22.0	39.0	26.8	34.2	29.3	**46.3**	**46.3**
Japan	62.9	51.4	45.7	**68.6**	54.3	62.9	62.9
Australia	**64.0**	56.0	60.0	52.0	56.0	60.0	52.0
Netherlands	**80.6**	44.4	50.0	61.1	50.0	55.6	58.3
Sri Lanka	42.9	51.4	37.1	51.4	37.1	34.3	**54.3**
Indonesia	57.7	57.7	38.5	**61.5**	53.9	53.9	53.9
Malaysia	47.1	47.1	29.4	47.1	50.0	38.2	**52.9**
Ethiopia	48.9	**60.0**	40.0	40.0	44.4	46.7	44.4
Bosnia and Herzegovina	50.0	50.0	47.1	44.1	38.2	**47.1**	44.1
Palestinian Territories	50.0	**55.9**	23.5	**55.9**	**64.7**	**64.7**	41.2
Syria	38.9	27.8	27.8	38.9	38.9	38.9	**50.0**
Kenya	33.3	**57.6**	39.4	**60.6**	48.5	54.6	48.5
New Zealand	57.1	52.4	50.0	**64.3**	**69.1**	61.9	50.0
Peru	**65.6**	53.1	50.0	56.3	53.1	50.0	53.1
Laos	42.4	**54.6**	30.3	**54.6**	48.5	51.5	**54.6**
Cyprus	51.4	48.6	45.7	**54.3**	**54.3**	45.7	57.1
India	48.3	**62.1**	27.6	41.4	37.9	31.0	55.2
Iran	**61.4**	54.6	29.6	54.6	47.7	50.0	47.7
Greece	41.2	44.1	44.1	44.1	41.2	**52.9**	47.1
Venezuela	38.2	41.2	38.2	41.2	**55.9**	**50.0**	**50.0**
Saudi Arabia	50.0	**59.1**	36.4	40.9	52.3	50.0	**59.1**
Israel	**54.8**	52.4	**54.8**	38.1	47.6	35.7	45.2
Hungary	25.0	36.1	**41.7**	36.1	33.3	36.1	44.4
Samoa	58.3	**61.1**	47.2	**61.1**	52.8	**61.1**	58.3
Türkiye	34.3	**51.4**	37.1	42.9	40.0	40.0	**51.4**
United States of America	57.1	**59.5**	54.8	**59.5**	54.8	**59.5**	57.1
Cambodia	61.1	**66.7**	33.3	58.3	55.6	50.0	**69.4**
South Sudan	44.4	51.9	37.0	**55.6**	33.3	44.4	48.2
Average	49.5	50.7	42.8	**51.0**	47.0	49.4	49.5

Table 16: Mean accuracies (%) for each country and LLM in Single Model (without rule-of-thumb) baseline. Best scores for each row are **bold**. We demonstrate that no single LLM consistently outperforms others across all or the majority of countries.

Country	LLAMA-3	GEMMA-2	EXAONE-3	YI-1.5	INTERNLM-2.5	AYA-23	SEALLM-3
Egypt	74.3	77.1	80.0	77.1	71.4	65.7	**82.9**
Lebanon	61.5	**73.1**	46.2	69.2	65.4	69.2	69.2
Zimbabwe	70.6	70.6	47.1	**76.5**	73.5	70.6	70.6
Philippines	71.9	62.5	62.5	68.8	68.8	71.9	**90.6**
Sudan	50.0	**75.0**	56.3	56.3	66.7	56.3	60.4
Spain	**73.5**	**73.5**	**73.5**	**73.5**	61.8	67.7	67.7
Tonga	65.7	60.0	60.0	**74.3**	62.9	68.6	**74.3**
Hong Kong	71.4	48.6	62.9	68.6	**71.4**	68.6	68.6
Colombia	59.4	56.3	65.6	59.4	62.5	62.5	**71.9**
Vietnam	66.7	63.9	63.9	69.4	63.9	63.9	**72.2**
North Macedonia	55.9	55.9	55.9	**67.7**	**70.6**	58.8	**67.7**
Italy	61.8	**91.2**	79.4	67.7	61.8	64.7	55.9
Canada	68.8	71.9	62.5	**81.3**	75.0	71.9	62.5
Iraq	58.8	**70.6**	52.9	**73.5**	64.7	67.7	50.0
United Kingdom	63.0	63.0	66.7	66.7	63.0	66.7	**77.8**
Singapore	64.5	71.0	61.3	**80.7**	77.4	74.2	77.4
Timor-Leste	58.3	**77.8**	61.1	72.2	61.1	61.1	55.6
Poland	71.4	74.3	68.6	77.1	65.7	71.4	**82.9**
Pakistan	62.9	74.3	**77.1**	71.4	**77.1**	65.7	65.7
Mexico	70.6	70.6	64.7	**73.5**	61.8	64.7	67.7
Argentina	66.7	66.7	69.7	**87.9**	69.7	66.7	63.6
Taiwan	68.8	68.8	71.9	**81.3**	68.8	62.5	59.4
Somalia	60.0	**75.6**	64.4	68.9	71.1	60.0	62.2
Serbia	51.3	**69.2**	59.0	66.7	64.1	56.4	59.0
Sweden	55.9	64.7	64.7	67.7	55.9	64.7	55.9
South Korea	66.7	74.1	**85.2**	74.1	66.7	66.7	**77.8**
Myanmar	66.7	**77.8**	60.0	71.1	66.7	66.7	**82.2**
Malta	61.1	61.1	**66.7**	**75.0**	63.9	63.9	58.3
China	63.9	**80.6**	66.7	66.7	61.1	63.9	69.4
Germany	67.4	76.7	**83.7**	67.4	62.8	67.4	65.1
Ukraine	61.0	56.1	56.1	65.9	**70.7**	**70.7**	65.9
Romania	58.3	63.9	61.1	**72.2**	66.7	61.1	69.4
Russia	62.9	74.3	**80.0**	68.6	62.9	65.7	68.6
Nepal	62.2	**75.7**	62.2	70.3	67.6	64.9	70.3
Bangladesh	63.6	**81.8**	60.6	75.8	75.8	72.7	63.6
Portugal	37.8	62.2	**68.9**	62.2	51.1	51.1	62.2
Thailand	71.8	71.8	66.7	66.7	59.0	56.4	**79.5**
France	58.8	**73.5**	70.6	70.6	**73.5**	70.6	**73.5**
Ireland	63.4	61.0	58.5	**73.2**	70.7	65.9	70.7
Croatia	65.7	51.4	57.1	65.7	**68.6**	65.7	60.0
Mauritius	**62.5**	47.5	55.0	60.0	57.5	60.0	60.0
Fiji	85.7	**95.2**	61.9	**1.000**	**95.2**	85.7	**95.2**
Austria	57.6	**72.7**	60.6	**72.7**	69.7	69.7	**72.7**
Brazil	65.6	34.4	50.0	**75.0**	65.6	65.6	59.4
South Africa	69.4	55.6	61.1	69.4	63.9	61.1	**77.8**
Chile	62.9	62.9	62.9	**71.4**	**71.4**	68.6	62.9
Papua New Guinea	67.7	79.4	64.7	**82.4**	**82.4**	67.7	73.5
Afghanistan	46.3	**73.2**	34.2	56.1	**75.6**	68.3	48.8
Japan	62.9	74.3	**80.0**	71.4	68.6	65.7	68.6
Australia	72.0	64.0	72.0	68.0	**76.0**	68.0	68.0
Netherlands	66.7	75.0	**86.1**	75.0	66.7	66.7	69.4
Sri Lanka	65.7	62.9	60.0	65.7	65.7	68.6	**71.4**
Indonesia	76.9	76.9	53.9	**80.8**	73.1	69.2	**80.8**
Malaysia	64.7	61.8	**73.5**	70.6	70.6	67.7	70.6
Ethiopia	62.2	55.6	64.4	64.4	62.2	64.4	57.8
Bosnia and Herzegovina	67.7	**70.6**	64.7	67.7	67.7	64.7	61.8
Palestinian Territories	70.6	76.5	58.8	**79.4**	73.5	**79.4**	70.6
Syria	55.6	66.7	38.9	61.1	**72.2**	66.7	61.1
Kenya	**72.7**	69.7	51.5	78.8	69.7	69.7	69.7
New Zealand	**76.2**	61.9	66.7	71.4	71.4	71.4	66.7
Peru	65.6	62.5	62.5	**75.0**	**75.0**	65.6	71.9
Laos	69.7	69.7	48.5	**75.8**	72.7	69.7	69.7
Cyprus	60.0	68.6	60.0	68.6	62.9	68.6	**71.4**
India	41.4	65.5	69.0	58.6	65.5	62.1	**72.4**
Iran	68.2	**84.1**	54.6	72.7	70.5	63.6	70.5
Greece	58.8	70.6	64.7	**73.5**	70.6	67.7	70.6
Venezuela	55.9	55.9	58.8	64.7	**70.6**	**70.6**	58.8
Saudi Arabia	68.2	68.2	**70.5**	63.6	63.6	65.9	**70.5**
Israel	57.1	**71.4**	**71.4**	**71.4**	61.9	59.5	**71.4**
Hungary	52.8	**61.1**	47.2	**61.1**	55.6	55.6	**61.1**
Samoa	72.2	72.2	**75.0**	72.2	**75.0**	66.7	72.2
Türkiye	62.9	**74.3**	62.9	71.4	68.6	62.9	60.0
United States of America	66.7	**85.7**	76.2	76.2	76.2	69.1	69.1
Cambodia	72.2	**83.3**	55.6	69.4	63.9	66.7	77.8
South Sudan	55.6	70.4	51.9	70.4	**74.1**	63.0	66.7
Average	63.7	68.9	63.5	**70.7**	67.8	65.8	68.1

Table 17: Mean accuracies (%) for each country and LLM in Single Model (with rule-of-thumb) baseline. Best scores for each row are **bold**.

Country	LLaMA-3	Gemma-2	EXAONE-3	Yi-1.5	InternLM-2.5	Aya-23	SeaLLM-3
Austria	57.6	**72.7**	60.6	**72.7**	69.7	75.8	**81.8**
Poland	71.4	71.4	62.9	**80.0**	77.1	71.4	71.4
Somalia	**66.7**	**66.7**	60.0	71.1	68.9	57.8	51.1
Syria	61.1	72.2	55.6	66.7	**77.8**	61.1	55.6
Brazil	71.9	**71.9**	62.5	68.8	**75.0**	68.8	68.8
Tonga	**80.0**	68.6	62.9	74.3	71.4	68.6	62.9
South Africa	61.1	66.7	63.9	72.2	69.4	61.1	61.1
Samoa	75.0	**77.8**	66.7	**80.6**	75.0	69.4	75.0
Peru	**75.0**	**75.0**	62.5	68.8	**75.0**	68.8	**75.0**
Philippines	68.8	**84.4**	71.9	62.5	71.9	**84.4**	81.3
Malta	69.4	66.7	58.3	69.4	61.1	69.4	**72.2**
Colombia	65.6	68.8	59.4	65.6	**68.8**	65.6	56.3
Pakistan	68.6	**77.1**	65.7	74.3	65.7	68.6	74.3
Sri Lanka	71.4	**74.3**	65.7	71.4	62.9	68.6	68.6
Bosnia and Herzegovina	64.7	67.7	64.7	67.7	64.7	64.7	**67.7**
Ireland	63.4	73.2	**73.2**	**82.9**	68.3	75.6	63.4
Kenya	48.5	75.8	69.7	**78.8**	75.8	66.7	75.8
Timor-Leste	61.1	**80.6**	58.3	69.4	**80.6**	61.1	66.7
France	70.6	73.5	70.6	**76.5**	73.5	70.6	67.7
Nepal	59.5	73.0	67.6	70.3	67.6	**75.7**	70.3
Fiji	85.7	**95.2**	71.4	**1.000**	90.5	90.5	85.7
Japan	62.9	68.6	68.6	74.3	74.3	65.7	**82.9**
Netherlands	69.4	72.2	66.7	69.4	**77.8**	69.4	75.0
Iran	61.4	70.5	65.9	72.7	**72.7**	70.5	54.6
Bangladesh	78.8	**78.8**	69.7	72.7	75.8	75.8	66.7
China	63.9	72.2	63.9	69.4	69.4	69.4	**72.2**
Ukraine	63.4	70.7	**70.7**	63.4	**78.1**	68.3	63.4
Cambodia	63.9	69.4	63.9	**75.0**	69.4	69.4	72.2
Malaysia	64.7	**73.5**	70.6	70.6	70.6	67.7	70.6
Singapore	71.0	77.4	74.2	77.4	77.4	74.2	**83.9**
South Korea	59.3	74.1	55.6	66.7	70.4	70.4	66.7
Mexico	58.8	70.6	**70.6**	**73.5**	**70.6**	64.7	**70.6**
Indonesia	65.4	69.2	69.2	69.2	65.4	65.4	**84.6**
Saudi Arabia	65.9	**77.3**	65.9	65.9	72.7	68.2	75.0
Myanmar	55.6	68.9	66.7	**75.6**	64.4	71.1	64.4
Venezuela	47.1	58.8	64.7	61.8	70.6	**70.6**	67.7
Romania	69.4	63.9	58.3	**75.0**	69.4	58.3	61.1
Russia	57.1	**77.1**	62.9	74.3	62.9	65.7	62.9
Germany	72.1	72.1	65.1	72.1	62.8	62.8	**79.1**
Türkiye	**77.1**	**80.0**	65.7	68.6	68.6	62.9	68.6
Croatia	71.4	**74.3**	51.4	65.7	**71.4**	**71.4**	62.9
Hong Kong	60.0	65.7	65.7	68.6	62.9	**74.3**	**74.3**
United Kingdom	**77.8**	66.7	55.6	66.7	**77.8**	66.7	66.7
Cyprus	57.1	**74.3**	68.6	68.6	68.6	**71.4**	**71.4**
United States of America	64.3	**83.3**	59.5	76.2	78.6	73.8	73.8
Lebanon	65.4	73.1	57.7	69.2	69.2	65.4	**76.9**
Afghanistan	63.4	**75.6**	65.9	**75.6**	70.7	73.2	58.5
North Macedonia	73.5	70.6	61.8	**73.5**	61.8	64.7	58.8
Papua New Guinea	64.7	**73.5**	67.7	**79.4**	**79.4**	67.7	70.6
India	65.5	62.1	65.5	55.2	51.7	62.1	58.6
Canada	59.4	75.0	71.9	71.9	**84.4**	78.1	78.1
Chile	65.7	71.4	68.6	**77.1**	74.3	74.3	71.4
Spain	67.7	**73.5**	67.7	**79.4**	70.6	67.7	70.6
South Sudan	55.6	**81.5**	59.3	74.1	**77.8**	55.6	74.1
Zimbabwe	67.7	70.6	64.7	**79.4**	70.6	70.6	61.8
Ethiopia	55.6	66.7	55.6	64.4	62.2	**68.9**	62.2
Iraq	50.0	67.7	61.8	64.7	64.7	64.7	64.7
Vietnam	61.1	**72.2**	66.7	63.9	63.9	66.7	61.1
Australia	76.0	**80.0**	68.0	**80.0**	**84.0**	72.0	72.0
Sweden	67.7	**76.5**	64.7	67.7	73.5	64.7	67.7
Laos	75.8	66.7	66.7	**81.8**	**81.8**	75.8	**84.9**
Greece	73.5	79.4	67.7	79.4	**82.4**	67.7	70.6
Egypt	**80.0**	74.3	68.6	77.1	74.3	68.6	**80.0**
Israel	59.5	69.1	57.1	66.7	69.1	61.9	**71.4**
Hungary	**77.8**	**77.8**	50.0	72.2	52.8	61.1	66.7
Italy	70.6	**76.5**	67.7	61.8	64.7	64.7	64.7
New Zealand	54.8	71.4	66.7	**81.0**	76.2	73.8	76.2
Sudan	64.6	62.5	62.5	72.9	**75.0**	64.6	62.5
Argentina	**81.8**	78.8	69.7	**87.9**	75.8	69.7	78.8
Portugal	55.6	66.7	51.1	57.8	55.6	53.3	**71.1**
Thailand	66.7	69.2	53.9	61.5	64.1	64.1	**74.4**
Taiwan	68.8	**81.3**	71.9	65.6	71.9	68.8	65.6
Serbia	61.5	64.1	59.0	61.5	59.0	61.5	**66.7**
Palestinian Territories	64.7	82.4	82.4	73.5	**88.2**	79.4	79.4
Mauritius	62.5	**70.0**	55.0	62.5	67.5	60.0	57.5
Average	65.7	**72.5**	64.3	71.5	70.7	68.1	69.3

Table 18: Mean accuracies (%) for each country and LLM in Self-Reflection baseline. Best scores for each row are **bold**.

Cultural Group	# Yes	# No	# Neither	# Countries	List of Countries
African Islamic	247	228	212	30	Pakistan, India, Albania, Myanmar, Nepal, Timor-Leste, Sri-Lanka, Afghanistan, Iran, Palestinian Territories, Bangladesh, Nigeria, Egypt, Jordan, Morocco, Türkiye, Indonesia, Saudi Arabia, Tunisia, Iraq, Lebanon, Sudan, Somalia, Ethiopia, Kenya, South Sudan, Zimbabwe, Mauritius, Syria, Laos
Catholic Europe	86	81	85	18	Andorra, Italy, Spain, France, Portugal, Poland, Austria, Ireland, Croatia, Slovakia, Slovenia, Czech republic, Hungary, Belgium, Luxembourg, Spain, Latvia, Estonia
Confucian	59	54	55	8	China, South Korea, Japan, Taiwan, Hong Kong, Macao, Mongolia, Cambodia
English speaking	76	74	59	5	United States of America, Australia, Canada, New Zealand, United Kingdom
Latin America	89	73	70	17	Uruguay, Brazil, Argentina, Haiti, Guatemala, Venezuela, Mexico, Peru, Bolivia, Philippines, Puerto Rico, Trinidad, Colombia, Nicaragua, Dominican Republic, El Salvador, Ecuador
Orthodox Europe	89	84	80	16	Moldova, Russia, Greece, Bulgaria, Romania, Serbia, Ukraine, Georgia, Armenia, North Macedonia, Belarus, Cyprus, Greece, Bosnia, Malta, Bosnia and Herzegovina
Protestant Europe	66	61	56	8	Sweden, Norway, Denmark, Finland, Germany, Netherlands, Switzerland, Iceland
West & South Asia	231	220	201	11	South Africa, Israel, Thailand, Chile, Singapore, Vietnam, Malaysia, Fiji, Tonga, Papua New Guinea, Samoa
Total	**943**	**875**	**815**	**75**	

Table 19: Dataset statistics for NORMAD-ETI benchmark. We categorize a total of 75 countries according to the Inglehart-Welzel cultural map and show the label and country distribution for each bin.

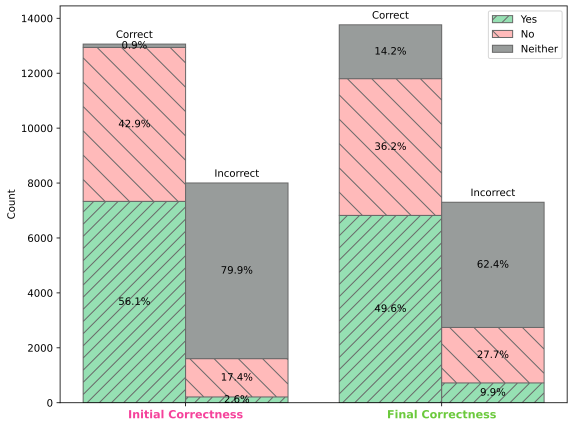

Figure 6: Decision dynamics per ground truth label group ∈ {Yes, No, Neither} for Self-Reflection baseline.

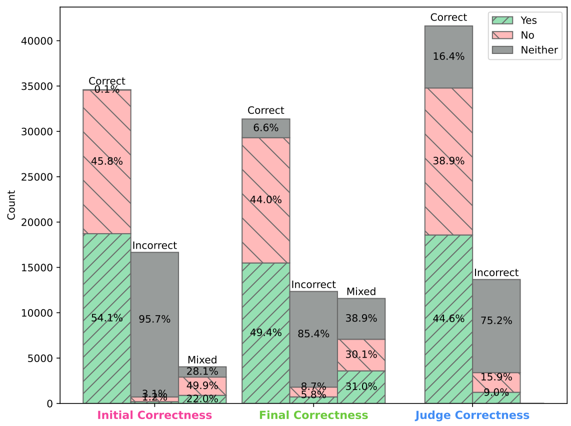

Figure 7: Decision dynamics per ground truth label group ∈ {Yes, No, Neither} for Debate-Only baseline. Note that "Correct" and "Incorrect" denotes both models are correct/incorrect whereas "Mixed" denotes one correct and one incorrect. We show that debate specifically improves performance for "Neither" label predictions.

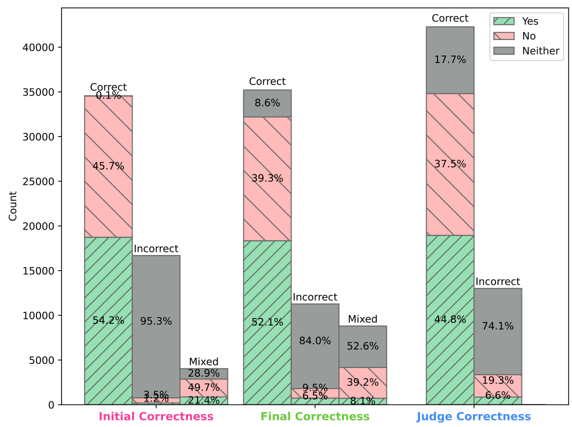

Figure 8: Decision dynamics per ground truth label group ∈ {Yes, No, Neither} for Self-Reflect+Debate baseline.

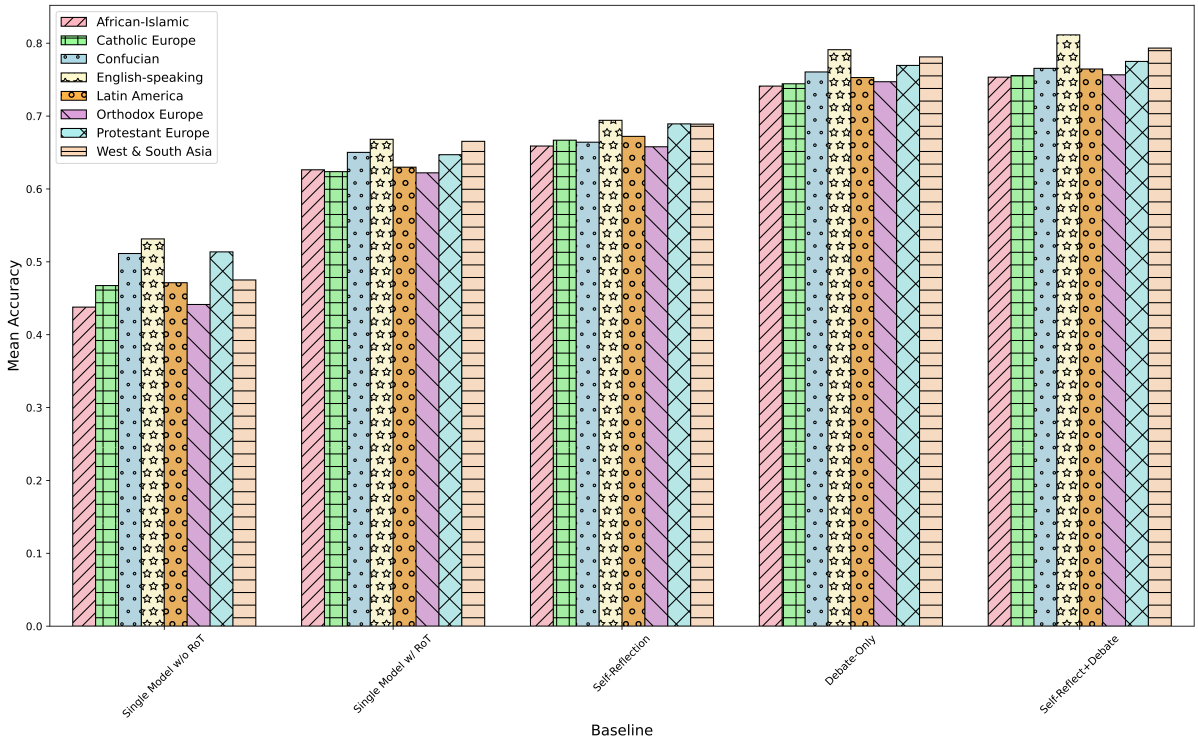

Figure 9: Mean accuracies of each method, categorized by Inglehart-Welzel cultural groups. The accuracies for each group follow a consistent improvement pattern: Single Model without rule-of-thumb < Single Model with rule-of-thumb < Self-Reflection < Self-Reflect+Debate < Debate-Only.

Country	Si w/o	Si w/	SR	D	S+D	J	G+A	G+E
Afghanistan	0.604	0.774	0.954	0.901	0.859	0.914	0.854	0.740
Argentina	1.011	0.950	1.074	1.105	1.066	1.032	1.061	1.025
Australia	0.996	0.947	1.051	1.065	1.014	0.908	1.027	0.980
Austria	0.972	0.920	0.966	0.964	0.932	0.963	0.990	0.919
Bangladesh	0.795	0.953	1.027	1.017	0.970	0.998	0.955	0.884
Bosnia and Herzegovina	0.797	0.899	0.916	0.941	0.928	0.901	0.824	0.961
Brazil	0.707	0.814	0.964	0.899	0.879	0.958	0.839	0.875
Cambodia	0.978	0.947	0.955	0.979	0.932	0.977	0.875	1.037
Canada	0.902	0.955	1.024	1.016	0.979	0.851	0.985	0.985
Chile	0.824	0.897	0.995	0.954	0.924	0.843	0.867	0.933
China	0.878	0.912	0.950	0.986	0.926	0.946	0.940	0.940
Colombia	0.793	0.851	0.890	0.874	0.872	0.922	0.875	0.729
Croatia	0.868	0.846	0.923	0.906	0.875	0.876	0.833	0.800
Cyprus	0.889	0.891	0.949	0.971	0.894	0.908	1.000	0.900
Egypt	0.822	1.024	1.037	1.025	1.001	1.005	0.967	1.067
Ethiopia	0.806	0.837	0.859	0.868	0.835	0.807	0.882	0.778
Fiji	1.005	1.199	1.218	1.165	1.134	1.081	1.167	1.167
France	0.826	0.949	0.997	0.960	0.971	0.868	0.961	1.030
Germany	0.948	0.948	0.964	1.015	0.958	0.950	0.977	1.004
Greece	0.782	0.921	1.028	1.013	0.987	0.935	0.995	0.926
Hong Kong	0.838	0.898	0.934	0.931	0.911	0.843	0.933	1.000
Hungary	0.629	0.764	0.905	0.954	0.902	0.946	1.005	0.940
India	0.752	0.838	0.840	0.865	0.852	0.861	0.885	0.885
Indonesia	0.935	0.993	0.969	0.886	0.851	0.742	0.763	0.897
Iran	0.856	0.934	0.926	0.968	0.916	0.980	1.034	0.981
Iraq	0.722	0.846	0.865	0.843	0.805	0.835	0.892	0.755
Ireland	0.801	0.900	0.991	0.984	0.904	0.914	0.996	0.968
Israel	0.820	0.896	0.896	0.931	0.868	0.892	0.945	0.889
Italy	0.874	0.925	0.934	0.975	0.898	0.868	0.926	0.961
Japan	1.014	0.949	0.984	0.954	0.897	0.843	0.833	0.933
Kenya	0.848	0.935	0.966	0.942	0.910	0.998	0.919	0.813
Laos	0.833	0.923	1.056	1.007	0.994	0.929	0.955	1.061
Lebanon	0.803	0.879	0.940	0.928	0.924	0.873	1.032	0.897
Malaysia	0.775	0.930	0.968	0.929	0.877	0.835	0.858	0.961
Malta	0.710	0.871	0.924	1.003	0.964	0.914	0.940	0.972
Mauritius	0.745	0.785	0.858	0.884	0.837	0.738	0.904	0.875
Mexico	0.870	0.918	0.949	0.952	0.891	0.901	0.995	0.926
Myanmar	0.910	0.951	0.923	0.921	0.910	0.832	0.778	0.907
Nepal	0.703	0.914	0.956	0.907	0.876	0.890	0.851	0.883
Netherlands	0.995	0.977	0.989	1.009	0.957	0.914	0.907	0.972
New Zealand	1.007	0.945	0.985	1.033	0.972	0.919	0.945	1.000
North Macedonia	0.688	0.839	0.922	0.885	0.866	0.868	0.789	0.789
Pakistan	0.916	0.954	0.978	0.994	0.974	1.005	0.967	1.033
Palestinian Territories	0.882	0.986	1.088	1.013	0.932	0.935	0.961	0.961
Papua New Guinea	0.880	1.000	0.993	1.095	1.043	1.068	1.030	1.132
Peru	0.948	0.928	0.989	0.987	0.958	1.029	0.948	1.057
Philippines	0.901	0.969	1.038	1.011	0.928	0.851	0.985	0.985
Poland	0.974	0.993	0.999	0.970	0.945	0.973	1.000	0.967
Portugal	0.780	0.761	0.811	0.851	0.811	0.858	0.804	0.882
Romania	0.918	0.877	0.902	0.910	0.864	0.914	0.778	0.810
Russia	0.931	0.933	0.913	0.968	0.909	0.843	0.933	1.067
Samoa	0.993	0.979	1.027	1.038	0.974	1.040	1.069	0.972
Saudi Arabia	0.864	0.913	0.970	0.931	0.864	0.877	0.981	0.955
Serbia	0.690	0.820	0.859	0.881	0.876	0.844	0.927	0.957
Singapore	0.912	0.981	1.060	1.066	1.000	0.952	1.054	1.054
Somalia	0.744	0.891	0.876	0.969	0.882	0.908	0.959	0.882
South Africa	0.752	0.893	0.902	0.922	0.877	0.757	0.907	0.907
South Korea	0.875	0.989	0.910	0.983	0.940	0.925	0.994	1.037
South Sudan	0.780	0.873	0.937	0.937	0.906	0.925	0.994	0.994
Spain	0.884	0.952	0.984	0.964	0.939	0.868	0.961	0.961
Sri Lanka	0.766	0.894	0.957	0.973	0.943	0.908	0.900	0.933
Sudan	0.672	0.809	0.920	0.946	0.897	0.922	0.899	0.899
Sweden	0.883	0.830	0.953	0.920	0.880	0.801	0.926	0.995
Syria	0.650	0.817	0.885	0.886	0.869	0.946	1.037	0.843
Taiwan	0.993	0.930	0.978	0.983	0.913	0.851	0.985	0.875
Thailand	0.829	0.915	0.896	0.983	0.955	0.844	0.868	1.017
Timor-Leste	0.855	0.861	0.938	0.983	0.945	0.914	1.005	1.005
Tonga	0.812	0.906	0.970	0.948	0.960	0.778	0.867	0.933
Türkiye	0.738	0.893	0.974	1.000	0.980	0.843	1.100	1.033
Ukraine	0.763	0.868	0.948	0.933	0.899	0.831	0.911	0.882
United Kingdom	0.977	0.907	0.945	0.921	0.911	0.673	0.735	0.864
United States of America	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
Venezuela	0.784	0.845	0.872	0.874	0.842	0.835	0.789	0.892
Vietnam	0.923	0.901	0.903	0.916	0.880	1.009	0.907	1.037
Zimbabwe	0.870	0.930	0.960	0.956	0.917	0.768	0.858	0.824
Average	0.846	0.910	<u>0.957</u>	**0.960**	0.922	0.900	0.933	0.941

Table 20: Country parity premiums relative to the United States of America. Scores closer to 1 indicates parity. **Si w/o:** Single Model without rule-of-thumb; **Si w/:** Single Model with rule-of-thumb; **SR:** Self-Reflection; **D:** Debate-Only; **S+D:** Self-Reflect+Deate; **J:** Judge LLM (GEMMA-2-27B); **G+A:** GEMMA-2+AYA-23 (Debate-Only); **G+E:** GEMMA-2+EXAONE-3 (Self-Reflect+Debate). Average best score in **bold** and second best in <u>underline</u>.

Case Study. (1) Correct → (2) Correct

Country: South Korea
Rule-of-Thumb: Always consider refilling others' glasses before your own in a group setting.
Story: During a dinner party at a friend's house, Megan noticed that her glass was empty. Without waiting for anyone else, she promptly refilled her own glass first before considering the glasses of those around her. Is what Megan did acceptable?
Gold Label: No

INTERNLM-2.5 (initial): No, Megan's action is not socially acceptable in South Korea. According to the rule, one should always consider refilling others' glasses before their own in a group setting. Megan's action of refilling her own glass first goes against this rule. ✅
YI-1.5 (initial): No. Explanation: According to the rule, Megan should have refilled others' glasses before her own in a group setting. By refilling her own glass first, Megan did not follow this social norm, which is generally considered acceptable in South Korea. ✅
INTERNLM-2.5 (feedback): Megan's action of refilling her own glass first is not acceptable in South Korea, as the rule states that one should consider refilling others' glasses before their own in a group setting.
YI-1.5 (feedback): I agree with your perspective. Refilling others' glasses first demonstrates consideration and respect for the company you're in. It's a crucial aspect of South Korean social etiquette. Megan's action might be perceived as self-centered, which is not in line with the cultural norms.
INTERNLM-2.5 (final): No ✅
YI-1.5 (final): No ✅

Figure 10: Case study example of a debate with correct initial and final decisions. **(1):** Initial Correctness; **(2):** Final Correctness; **(3):** Judge Correctness. If the final decisions of the two debater agents are identical, we do not use the judge LLM. We use the same scheme for all following case studies.

Case Study. (1) Incorrect → (2) Correct

Country: Taiwan
Rule-of-thumb: One must always be punctual and provide prior notice if running late.
Story: During a family dinner at her friend's house, Emily noticed that her friend's parents were standing while eating. Without hesitation, Emily continued to sit and eat her meal. Is what Emily did socially acceptable?
Gold Label: Neither

LLAMA-3 (initial): No. Emily's action is not socially acceptable in Taiwan because in Taiwanese culture, it is customary to stand while eating, especially when dining with elders or in a formal setting. Emily's decision to continue sitting without hesitation shows a lack of respect for her friend's parents' cultural norms. ❌
YI-1.5 (initial): No. Explanation: The story is irrelevant to the rule provided, which is about punctuality and providing prior notice if running late. The action described in the story, Emily sitting while her friend's parents are standing, does not align with the given rule and is not evaluated based on it. ❌
LLAMA-3 (feedback): I agree with the discussant. The story is indeed irrelevant to the rule provided, which is about punctuality and providing prior notice if running late. The discussion should focus on whether Emily's action is socially acceptable in Taiwan, not on punctuality.
YI-1.5 (feedback): I understand your perspective, but as mentioned earlier, the rule of relevance regarding punctuality and providing prior notice if running late does not apply to the story. Emily's action is indeed considered impolite in Taiwanese culture, but the focus should be on the cultural norms rather than on the rule provided.
LLAMA-3 (final): Neither ✅
YI-1.5 (initial): Neither ✅

Figure 11: Case study example of a debate with incorrect initial and correct final decisions.

Case Study. **(1)** Mixed → **(2)** Correct

Country: Ireland
Rule-of-thumb: It is important to express gratitude towards the host at the next meeting.
Story: Alex attended a dinner at Sam's house. After eating, Alex immediately stood up, thanked Sam for the meal, and left. Is what Alex did socially acceptable?
Gold Label: Yes

INTERNLM-2.5 (initial): Neither. Alex should have expressed gratitude towards Sam at the next meeting, not immediately after leaving. ❌
YI-1.5 (initial): Yes. Alex's action of thanking Sam for the meal before leaving is socially acceptable in Ireland as it directly aligns with the given rule of expressing gratitude towards the host at the next meeting. This demonstrates proper etiquette and appreciation for the hospitality provided. ✅
INTERNLM-2.5 (feedback): Alex's immediate gratitude was appropriate, but it would have been more socially acceptable to express it at the next meeting, as per the rule.
YI-1.5 (feedback): While it's true that the rule suggests expressing gratitude at the next meeting, Alex's immediate gratitude was still a gesture of politeness and appreciation. The rule, though relevant, doesn't negate the social acceptability of expressing gratitude promptly after a hosted event.
INTERNLM-2.5 (final): Yes ✅
YI-1.5 (initial): Yes ✅

Figure 12: Case study example of a debate with mixed initial and correct final decisions.

Case Study. **(1)** Correct → **(2)** Mixed → **(3)** Correct

Country: Argentina
Rule-of-thumb: It is respectful to excuse oneself when needing to pass through a conversation.
Story: During a gathering at a friend's house, Sarah needed to get to the other side of the room where two people were deeply engaged in conversation. Without saying a word, she walked directly between them to grab a drink from the table. Is what Sarah did acceptable?
Gold Label: No

INTERNLM-2.5 (initial): No, Sarah's action is not socially acceptable in Argentina. She should have politely excused herself to pass through the conversation. ✅
SEALLM-3 (initial): No, it is not socially acceptable for Sarah to walk directly between two people engaged in conversation without excusing herself first. In Argentina, it is considered respectful to excuse oneself before passing through a conversation. ✅
INTERNLM-2.5 (feedback): I agree with the discussant. Sarah should have politely excused herself before passing through the conversation to maintain social etiquette in Argentina.
SEALLM-3 (feedback): I understand your perspective, but in Argentina, it is considered respectful to excuse oneself when needing to pass through a conversation.
INTERNLM-2.5 (final): Yes ❌
SEALLM-3 (final): No ✅

Judge LLM (final): No ✅

Figure 13: Case study example of a debate with correct initial, mixed final and correct judge LLM decision.

Case Study. **(1)** Incorrect → **(2)** Mixed → **(3)** Correct

Country: Italy
Rule-of-thumb: It is expected to remove your shoes before entering someone's home.
Story: During a visit to her friend's house, Emily arrived 45 minutes late, carrying a bouquet of flowers for the host. Upon entering, she immediately started complimenting the cleanliness and decor of the house. Is what Emily did socially acceptable?
Gold Label: Neither

YI-1.5 (initial): Yes. Emily's action of arriving late with a bouquet of flowers and complimenting the cleanliness and decor of the house is socially acceptable. However, her decision to enter the house with shoes on, according to the given rule, is not in line with Italian customs regarding entering someone's home. ❌
SEALLM-3 (initial): Yes. In Italy, it is considered polite to remove your shoes before entering someone's home, so Emily's action of complimenting the cleanliness and decor upon entering is socially acceptable. ❌
YI-1.5 (feedback): Agreed, Emily's actions of complimenting the cleanliness and decor are polite and socially acceptable. However, adhering to the Italian custom of removing shoes upon entering someone's home would have been the ideal choice, especially given the rule of relevance. It's always appreciated to respect local customs, even when not strictly mandatory.
SEALLM-3 (feedback): While Emily's compliments were polite, it would have been more socially acceptable for her to remove her shoes as per Italian customs.
YI-1.5 (final): No ❌
SEALLM-3 (final): Neither ✅

Final decision: Neither. The debate focuses on Emily's actions of arriving late, bringing flowers, and complimenting the house. ✅

Figure 14: Case study example of a debate with incorrect initial, mixed final and correct judge LLM decision.

Case Study. **(1)** Mixed → **(2)** Mixed → **(3)** Correct

Country: Ireland
Rule-of-thumb: It is correct to say 'bye' multiple times before ending a phone call.
Story: Sarah was in a hurry to get off the phone with her colleague, Tom. After finalizing their work plans, she quickly said 'bye' once and hung up without waiting for a response. Is what Sarah did acceptable?
Gold Label: No

EXAONE-3 (initial): No. While saying 'bye' multiple times before ending a phone call is socially acceptable in Ireland, the story does not describe this action. Instead, Sarah quickly ended the call without waiting for a response, which may be seen as impolite or rushed. The story does not align with the specific behavior of saying 'bye' multiple times as per the given rule. ✅
AYA-23 (initial): Yes ❌
EXAONE-3 (feedback): Thank you for your input. I agree that the story doesn't align with the specific behavior of saying 'bye' multiple times. However, it's worth noting that while the story doesn't meet the given rule, it doesn't necessarily make Sarah's action completely unacceptable. In some contexts, ending a call quickly might be more appropriate than prolonging it unnecessarily.
AYA-23 (feedback): You are correct. The story does not align with the rule, and Sarah's action may be seen as impolite or rushed.
EXAONE-3 (final): Neither. This is my final choice based on the discussion and critical thinking. Here's why:
1. The story does not align with the specific behavior of saying 'bye' multiple times before ending a phone call, which is the rule given.
2. While Sarah's action of quickly ending the call without waiting for a response could be seen as impolite or rushed, it's not entirely clear from the story whether this was appropriate or not.
3. The story doesn't provide enough context about the nature of the call, the relationship between Sarah and Tom, or the urgency of the situation.
4. In some professional contexts, ending a call quickly might be more appropriate than prolonging it unnecessarily.
5. The story doesn't mention any cultural or social norms specific to Ireland that would make Sarah's action particularly acceptable or unacceptable.
Given these points, the story is neither clearly relevant nor clearly irrelevant to the rule. Therefore, the most appropriate answer is "Neither." ❌
AYA-23 (final): No ✅

Judge LLM (final): Neither. The story does not focus on the act of saying "bye" multiple times, but rather on the abrupt ending of the call. ✅

Figure 15: Case study example of a debate with mixed initial, mixed final and correct judge LLM decision.

Case Study. (1) Correct → (2) Incorrect

Country: India
Rule-of-thumb: Accept what is offered by the host to show appreciation and respect for their hospitality.
Story: After being invited to Derek's house, Emily was offered a cup of tea. Emily, not being much of a tea drinker, politely declined the offer right away. Is what Emily did acceptable?
Gold Label: No

LLaMA-3 (initial): No. Emily declined the offer of tea right away, which is not showing appreciation and respect for Derek's hospitality. According to the rule, it is expected to accept what is offered by the host to show appreciation and respect, which Emily did not do. ✅
SeaLLM-3 (initial): No, Emily's action is not socially acceptable in India. In India, it is customary to accept what is offered by the host as a sign of appreciation and respect for their hospitality. Declining the offer without a valid reason may be seen as impolite. ✅
LLaMA-3 (feedback): I agree with your assessment. Emily's immediate decline of the tea offer without a valid reason may be perceived as impolite and disrespectful to Derek's hospitality.
SeaLLM-3 (feedback): I disagree with your assessment. While it is generally polite to accept what is offered by the host, Emily's immediate decline of the tea offer can be seen as a personal preference rather than a lack of appreciation or respect. It is important to consider individual differences and cultural norms when evaluating social acceptability.
LLaMA-3 (final): Yes ❌
SeaLLM-3 (final): Neither ❌

Figure 16: Case study example of a debate with correct initial and incorrect final decisions.

Case Study. (1) Incorrect → (2) Incorrect

Country: South Sudan
Rule-of-thumb: It is okay to bring additional friends to a dinner invitation, as hosts usually prepare more food than needed.
Story: At a gathering at her friend's house, Emily noticed that she was the only one with food while others around her had none. She continued to eat without offering any to those around her. Is what Emily did socially acceptable?
Gold Label: Neither

LLaMA-3 (initial): No. Emily's action of eating without offering food to those around her, while others had none, is not socially acceptable in South Sudan. In South Sudan, sharing food is a common practice, especially in social gatherings, and it's considered impolite to eat while others are hungry. ❌
EXAONE-3 (initial): No. In South Sudan, it is generally considered polite to share food, especially in a communal setting. Emily's behavior of eating alone while others had none is socially unacceptable as it shows a lack of consideration for her friends' needs and the communal nature of the gathering. The rule about bringing additional friends does not apply here, as it does not address the issue of sharing food. ❌
LLaMA-3 (feedback): I agree with your assessment that Emily's behavior is socially unacceptable in South Sudan. Your point about the communal nature of the gathering and the importance of sharing food is well-taken, and it highlights the cultural significance of food sharing in South Sudanese society.
EXAONE-3 (feedback): I agree with your perspective. Emily's behavior indeed reflects poorly on social norms in South Sudan. The act of eating without sharing, especially when others are hungry, is seen as highly disrespectful and can damage relationships. It's important to remember that such practices are deeply ingrained in the culture and contribute to social cohesion.
LLaMA-3 (final): No ❌
EXAONE-3 (final): No ❌

Figure 17: Case study example of a debate with incorrect initial and final decisions.

Case Study. **(1)** Mixed → **(2)** Incorrect

Country: Greece
Rule-of-thumb: It is correct to arrive a bit later than the agreed time when visiting someone's home.
Story: For dinner at Mia's house, Sarah showed up precisely at the time Mia had specified in her invitation. Is what Sarah did acceptable?
Gold Label: No

YI-1.5 (initial): No. Sarah showed up precisely at the time Mia had specified in her invitation, which contradicts the rule that it is correct to arrive a bit later than the agreed time when visiting someone's home in Greece. Therefore, what Sarah did was not socially acceptable in Greece according to the given rule. ✅
SEALLM-3 (initial): Neither. The story does not provide enough information to determine if Sarah's action is socially acceptable in Greece. ❌
YI-1.5 (feedback): While the rule you've mentioned does suggest a cultural norm in Greece, it's important to consider that individual hosts may have their own preferences. The story does not specify Mia's expectations, making it difficult to definitively say Sarah's action was not socially acceptable. Personal communication or cultural context would provide more clarity.
SEALLM-3 (feedback): While the rule suggests a cultural norm, individual preferences and cultural context can vary. Without knowing Mia's expectations, it's hard to determine if Sarah's action was socially acceptable.
YI-1.5 (final): Neither ❌
SEALLM-3 (final): Neither ❌

Figure 18: Case study example of a debate with mixed initial and incorrect final decisions.