

CCD-Bench: Probing Cultural Conflict in Large Language Model Decision-Making

Hasibur Rahman, Hanan Salam

Center of Interdisciplinary Data Science & AI (CIDSAI), NYUAD Research Institute, New York University Abu Dhabi, Abu Dhabi, United Arab Emirates

hasibur.rahman@nyu.edu, hanan.salam@nyu.edu

Abstract

Large language models (LLMs) increasingly shape interpersonal and societal decision-making, yet their ability to navigate explicit conflicts between legitimate cultural values remains underexplored. Existing benchmarks focus on cultural knowledge (CulturalBench), value inference (World-ValuesBench), or single-axis bias (CDEval), but none assess how LLMs adjudicate when multiple cultural frameworks directly clash. We introduce **CCD-Bench (Culture-Conflict Decision Benchmark)**, a benchmark for evaluating LLM decision-making under cross-cultural value conflict. CCD-Bench contains 2,182 open-ended dilemmas across seven domains, each with ten anonymized response options aligned with the ten GLOBE cultural clusters spanning 62 societies. Using a Stratified Latin Square design, we evaluate 17 leading LLMs and find clear biases: models favor Nordic Europe (20.2%) and Germanic Europe (12.4%), while Eastern Europe and Middle East & North Africa responses are least preferred ($\approx 5\text{--}6\%$). Although 87.9% of model rationales reference multiple cultural dimensions, this pluralism is shallow, dominated by Future and Performance Orientation, with limited attention to Assertiveness or Gender Egalitarianism ($<3\%$). Ordering effects are negligible, and model similarity clusters by developer lineage rather than geography. CCD-Bench shifts evaluation from bias detection to pluralistic reasoning, revealing that current LLMs express a Western-centric, consensus-oriented worldview even when confronted with equally valid, culturally diverse alternatives.

Code — <https://github.com/smartlab-nyu/CCD-Bench>

Introduction

Large Language models (LLMs) are increasingly deployed as mediators in interpersonal and societal decision-making, including conflict resolution and moral judgment tasks (Tan et al. 2024; Pires et al. 2025; Koçak, Puranam, and Yeğın 2024). As conversational interfaces integrate LLMs into everyday applications (Lalwani, Saleh, and Salam 2025; Elgarf, Salam, and Peters 2024), millions of users now rely on them for socially situated guidance. Even a mundane query, such as “How should I email my boss to request leave?”, requires not only factual accuracy but also culturally sensitive judgment that aligns with local values and expectations.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

However, the global adoption of LLMs reveals a critical challenge: aligning model recommendations with the full spectrum of cultural value systems. A growing body of evidence shows that LLMs inherit salient cultural leanings from their training data (Wang et al. 2024; Karinshak et al. 2024; Lu, Song, and Zhang 2025; AlKhamissi et al. 2024), leading users to receive responses that, while coherent, may reflect narrow or inconsistent cultural assumptions. For instance, the question “How should responsibilities be shared within a family?” can evoke equally legitimate but divergent answers, from merit-based distribution to age-based hierarchy to gender egalitarianism, depending on cultural context. Although prompt-engineering strategies that foreground local contextual cues can mitigate such effects (Xu et al. 2025; Zhong, Yun, and Sun 2024), a deeper issue remains: “*How do LLMs resolve situations in which legitimate cultural values, such as hierarchy versus egalitarianism, or group harmony versus personal initiative, directly conflict?*”. These systematic preferences raise an important distinction between bias and cultural trend. Following AlKhamissi et al. (2024), we refer to these patterns as cultural trends rather than biases, emphasizing that they do not necessarily entail harm but highlight limitations in addressing pluralistic value systems.

Existing leading benchmarks, such as MMLU (Hendrycks et al. 2021) and ELO (Zheng et al. 2023), prioritize factual accuracy and stylistic quality, overlooking the cultural nuance and social intelligence required in real-world interactions. Early alignment paradigms similarly prioritize universal principles such as helpfulness, honesty, and harmlessness (the “HHH” criteria) (Askill et al. 2021; Ouyang et al. 2022; Bai et al. 2022), while sidestepping the pluralistic and culturally contingent nature of human values (Barbara Herrnstein Smith 1991). Yet, real decisions often hinge on reconciling legitimate but conflicting value systems, a challenge where the “HHH” framework offers little guidance. Existing cultural benchmarks (Wang et al. 2024; Chiu et al. 2024; Zhao et al. 2024) capture model tendencies along individual cultural axes, but none examine how LLMs reason when culturally grounded values directly collide. This leaves a critical gap in evaluating LLM decision-making within pluralistic, cross-cultural contexts.

We introduce **CCD-Bench (Culture-Conflict Decision Benchmark)**, a benchmark for evaluating how LLMs navigate value-plural dilemmas: cases where multiple cultural

worldviews are equally defensible. It assesses models’ ability to engage in pluralistic reasoning, recognizing and justifying choices across competing cultural values.

Our diagnostic suite evaluates the 17 widely used LLMs (excluding specialized reasoning models), embedding all culturally salient information directly in the prompt to isolate the LLMs’ intrinsic decision-making tendencies. The benchmark comprises 2,182 dilemmas across seven domains: Arts, Education, Family, Wellness, Work, Science, and Lifestyle, adapted from Wang et al. (2024). Building on evidence that many LLMs already encode cultural frameworks (Karinshak et al. 2024), we generate ten anonymized candidate answers for each dilemma using a state-of-the-art reasoning model, followed by human verification. Each response corresponds to one of the ten GLOBE cultural clusters: Anglo, Latin Europe, Nordic Europe, Germanic Europe, Eastern Europe, Latin America, Sub-Saharan Africa, Middle East & North Africa, Southern Asia, and Confucian Asia, representing 62 societies (House et al. 2004). Instead of locating models along isolated cultural dimensions, CCD-Bench embeds the cultural trade-off within each dilemma: the model must select and justify one of ten equally valid, culturally grounded responses. This design directly exposes how LLMs reason when confronted with competing cultural imperatives.

We advance three main contributions:

1. **Benchmark:** We introduce CCD-Bench, the first benchmark for evaluating LLMs on navigating complex cross-cultural value conflicts, comprising 2,182 dilemmas across seven domains, each paired with ten anonymized culturally-distinct responses representing the GLOBE cultural clusters.
2. **Methodology:** We develop a generation and verification pipeline for producing culturally authentic alternatives, combining multi-cluster prompting, similarity pruning, human annotation, and a Stratified Latin Square design to ensure a fair assessment of value pluralism.
3. **Analysis:** We evaluate 17 leading LLMs on CCD-Bench, revealing systematic patterns in how models navigate cultural conflicts and the extent to which they default to specific cultural frameworks.

Related Work

Cultural Alignment in LLMs. Compelling evidence shows LLMs encode distinct cultural orientations. For instance, prompting GPT-4 in English tends to produce analytic, independence-oriented reasoning, whereas Chinese prompts evoke holistic, interdependent values (Lu, Song, and Zhang 2025). Similarly, *CDEval* finds GPT-4 skewed toward low power distance and high individualism, while Baichuan-13B and Qwen-7B exhibit stronger collectivist tendencies (Wang et al. 2024). Yet, such binary, one-dimension-at-a-time evaluations oversimplify pluralism. *LLM-GLOBE* advances this direction using open-ended prompts to uncover tonal and rhetorical differences across cultures (Karinshak et al. 2024), while Alkhamissi et al. (2024) improves alignment through “Anthropological Prompting.” However, these

studies stop short of testing how models reason when legitimate cultural values directly conflict.

Cultural Knowledge and bias benchmarks. Several Benchmarks assess cultural understanding or bias in LLMs but similarly avoid explicit value conflict. For instance, *CulturalBench* (Chiu et al. 2024) features 1,227 human-written questions across 45 regions and 17 topics to assess cultural knowledge, with variants for different difficulty levels. However, it focuses on recall and ambiguity handling rather than decision-making amid conflicting norms. *World-ValuesBench* (Zhao et al. 2024) draws from the World Values Survey ¹ to create millions of examples for predicting demographic-conditioned values, evaluating alignment via distribution metrics, but emphasizes prediction over resolution of conflicting perspectives. *BLENd* (Myung et al. 2024) tests everyday cultural knowledge with 52,600 QA pairs across 16 countries and 13 languages, revealing performance disparities for low-resource cultures, although it targets factual recall rather than value trade-offs. *CAMeL* (Naous et al. 2024) and its extensions detect biases in associations, particularly between Arab and Western entities, using fill-in-the-blank prompts and parallel data, but center on implicit biases instead of explicit dilemmas requiring justification. These benchmarks illuminate cultural gaps but center on representation and factual recall rather than decision-making under normative conflict.

Despite significant progress, no existing work systematically examines how LLMs adjudicate cultural value conflicts: situations where multiple legitimate worldviews demand incompatible choices. Addressing this gap, CCD-Bench focuses on pluralistic dilemmas that require explicit reasoning about competing cultural imperatives.

Methodology

This section presents the methodology underlying CCD-Bench, a benchmark designed to evaluate how LLMs navigate culturally grounded value conflicts. We used GLOBE as it aligns with our dilemma structure, provides validated cultural value information for 62 societies, and distinguishes between practices and values (House et al. 2004). The 10-cluster granularity is sufficient for demographic analyses. The benchmark construction process follows a three-module pipeline encompassing Generation, Verification, and Evaluation. The Generation Module produces culturally diverse dilemmas and candidate responses aligned with the GLOBE framework; the Verification Module ensures authenticity and cross-cultural validity through human annotation and quality control; and the Evaluation Module standardizes presentation and model assessment procedures. Together, these modules establish a replicable framework for systematically examining LLM decision-making under cultural pluralism.

Generation Module

The Generation Module constructs the core of CCD-Bench, producing a diverse set of 2,182 open-ended cultural dilemmas, each paired with ten culturally grounded response alternatives. We build on the validated question corpus from

¹<https://www.worldvaluessurvey.org/wvs.jsp>

CDEval (Wang et al. 2024), the most comprehensive collection of culturally relevant scenarios to date, to ensure that CCD-Bench captures authentic cultural decision points while remaining comparable with prior cultural-AI evaluations. Starting from 2,953 short-form questions, we retain 2,204 open-ended items likely to elicit divergent advice across cultural contexts, excluding binary or factual questions. Each retained scenario is categorized into one of seven domains: Arts, Education, Family, Wellness, Work, Science, and Lifestyle, to represent everyday value conflicts.

Each dilemma is designed to contrast two or more competing cultural imperatives (e.g., hierarchy vs. egalitarianism, collectivism vs. individual autonomy). For each dilemma q , we generate a set of 10 *culturally distinctive answers* $\{a_c\}_{c=1}^{10}$, each corresponding to one of the ten GLOBE clusters $\mathcal{C} = \{\text{Anglo, Latin Europe, \dots, Confucian Asia}\}$ (House et al. 2004). To produce these alternatives, we employ a multi-cluster prompting strategy grounded in the GLOBE framework, instructing a state-of-the-art reasoning model (OpenAI o3-2025-04-16) to generate one response per cluster in a single call. Each prompt embeds fenced `<cultural identity>` directives that encapsulate anonymized GLOBE descriptors. The model is required to return: (i) a *one-sentence answer* that omits explicit cultural self-reference and (ii) a *one-sentence rationale* that references the corresponding GLOBE cultural dimensions. Following Balepur, Rudinger, and Boyd-Graber (2025), rationales are included to enhance internal coherence between cultural identity and decision justification. A detailed prompt structure is provided on GitHub.

Similarity pruning and regeneration. Although the multi-cluster prompting procedure should yield culturally diverse responses, we observed occasional lexical redundancy when answers from different clusters conveyed similar advice using overlapping phrasing. To mitigate this, we computed *pairwise MinHash Jaccard similarity* (Cann et al. 2025) over 3-gram shingles for each ten-tuple of responses. Any pair exceeding a similarity threshold of $\tau > 0.70$ triggered regeneration of the more redundant answer, with a maximum of two attempts. The threshold was determined empirically: manual inspection confirmed that pairs above 0.70 exhibited minimal semantic differentiation, whereas pairs below retained meaningful cultural distinctions. After pruning, 22 questions were removed, yielding 2,182 final dilemmas. Across these, the pairwise MinHash Jaccard similarity among the ten alternatives had a mean of $M = 0.123$ with a standard deviation of $\sigma = 0.109$ ($N = \binom{10}{2} \times 2182$ comparisons), confirming that regeneration effectively reduced near-duplicate content while preserving linguistic and semantic diversity. This process enhanced the discriminative validity of CCD-Bench by ensuring that each response remains lexically distinct yet culturally coherent, a result further verified through subsequent human evaluation.

Option ordering. Pezeshkpour and Hruschka (2024) demonstrated that the order of answer options in multiple-choice settings can significantly affect LLM performance. To neutralize such serial-position effects, we employ a 10×10 *Stratified Latin Square*, a domain-specific adap-

tation of the standard Latin square approach. Independent Latin squares are generated for each of the seven domains to ensure balanced option ordering within domains, rather than globally across the corpus. Each cultural alternative appears exactly once in every ordinal position (1–10) and precedes or follows every other alternative exactly once, thereby satisfying both row and column orthogonality. The complete evaluation-prompt specification is available on GitHub to support transparency and reproducibility.

Verification Module

A team of five reviewers and one meta-reviewer confirmed that (i) each generated option would be viewed as *reasonable* within the target cluster, (ii) did not mention the cultural cluster directly, and (iii) was semantically distinct from the other nine options. Non-conforming options are regenerated until all criteria are met, with a maximum of two attempts. None of the earlier 2,182 dilemmas were removed at this step. More details can be found on GitHub.

Evaluation Module

Our analytic pipeline decomposes evaluation logs into three layers: descriptive statistics, null-hypothesis significance tests, and robust similarity/diversity indices.

Descriptive layer. For each evaluation file, we record (i) absolute frequency f_c of each cultural cluster $c \in \{1, \dots, 10\}$, (ii) display position distribution $\pi = (\pi_1, \dots, \pi_{10})$, and (iii) domain-specific cluster counts. Raw tallies are used without smoothing.

Position-bias test. Under null hypothesis H_0 of equal likelihood for $k = 10$ positions, let \mathbf{O} be observed counts and $\mathbf{E} = (N/k, \dots, N/k)$ expected counts, with $N = \sum_i O_i$. The χ^2 goodness-of-fit test yields

$$\chi^2 = \sum_{i=1}^k \frac{(O_i - E_i)^2}{E_i}, \quad V = \sqrt{\frac{\chi^2}{N(k-1)}}, \quad (1)$$

where Cramér’s V follows thresholds: negligible < 0.10 , small < 0.30 , medium < 0.50 , large ≥ 0.50 (Harald Cramér 1999).

Domain-cluster association. We construct a $D \times C$ contingency table $\mathbf{T} = [T_{ij}]$ for domains (D) vs. clusters (C). Pearson’s $\chi^2(\mathbf{T})$ with $(D-1)(C-1)$ degrees of freedom is reported with standardized residuals $r_{ij} = \frac{T_{ij} - E_{ij}}{\sqrt{E_{ij}}}$ to identify over/under-represented pairs (Agresti 2019).

Latin-square balance. Option orderings use a 10×10 Latin square (Bradley 1958). Let A be the 10×10 matrix of cluster appearances per position across Q questions. Imbalance is measured as

$$\Delta = \|A - (Q/10)\mathbf{1}\|_1, \quad \text{BS} = 100 \left(1 - \frac{\Delta}{1.8 \cdot 10Q} \right) \%, \quad (2)$$

with 1.8 as the worst-case normalization factor.

Diversity metrics. For cluster frequencies \mathbf{p} and sorted counts \mathbf{x} with $x_{(1)} \leq \dots \leq x_{(n)}$,




Figure 1: Percentage distribution of cultural cluster selections across 17 LLMs. Each marker represents one model, with values normalized to percentages (0–100%). The x-axis is limited to 35% because no model exceeds this value.

$$H(\mathbf{p}) = - \sum_c p_c \log_2 p_c, \quad (3)$$

$$G(\mathbf{x}) = \frac{1}{n} \left(n + 1 - \frac{2}{\sum_i x(i)} \sum_{i=1}^n (n + 1 - i) x(i) \right) \quad (4)$$

(Shannon and Weaver 1949; Dorfman 1979). For each statistic $\hat{\theta}$, we draw $B = 5,000$ bootstrap replicates and report [2.5, 97.5] percentiles as 95% confidence bands (Efron and Tibshirani 1998).

Model similarity. For LLMs with distributions P and Q , symmetrized Kullback–Leibler divergence (Kullback 1997):

$$D_{\text{sym}}(P, Q) = \frac{1}{2} \sum_c \left[P_c \log \frac{P_c}{Q_c} + Q_c \log \frac{Q_c}{P_c} \right], \quad (5)$$

adding 10^{-10} to avoid $\log 0$. The matrix feeds average-linkage hierarchical clustering (Sneath and Sokal 1973).

Cluster–position bias. Repeating position-bias per cluster gives Cramér’s V_c ; mean $\bar{V} = \frac{1}{10} \sum_c V_c$ is the model-bias score.

Pluralism in rationales. Let $I_i = 1$ if rationale i cites ≥ 2 GLOBE dimensions ($I_i = 0$ otherwise), for $i = 1, \dots, n$. The pluralism rate is $R_{\text{plural}} = \frac{1}{n} \sum_i I_i$, reported per LLM and ensemble to assess value-reasoning breadth.

Experiments

We experiment with the 17 most popular LLMs for non-technical tasks, as ranked by OpenRouter.² Our inquiry is structured around five research questions (RQs):

²<https://openrouter.ai/rankings>

When offered culturally valid options, do LLMs still prefer certain clusters or values? (RQ1)

A pronounced Nordic trend. Every single model in our sample selects the **Nordic Europe** option most often. The share ranges from 16.1% (Cohere `command-r`) to a remarkable 30.6% (Anthropic Claude–4 Sonnet), with an across-model mean of $M = 20.2\%$. Put differently, the median model chooses Nordic-style advice roughly $\times 2$ more often than chance, signaling an entrenched preference for highly egalitarian, consensus-driven solutions. Figure 1 shows the percentage selection frequencies for all 17 LLMs.

Secondary leanings. Two clusters emerge as consistent runners-up: **Germanic Europe** ($M = 12.4\%$): 7 LLMs list it as their second-most frequent choice. And **Sub-Saharan Africa** ($M = 11.5\%$): likewise second for seven LLMs, underscoring that collectivist but low-power-distance reasoning resonates with many systems. A smaller subset (three LLMs, all OpenAI or Google) places the **Anglo** cluster in second position, yet its average share (11.3%) still lags behind the Nordic and Germanic means.

Under-represented clusters. At the lower end of the spectrum lie **Eastern Europe** (5.6%) and the **Middle East & North Africa** (5.8%). No LLM exceeds an 8.2% selection rate for either group, and the minimum for Eastern Europe drops to just 2.7% in Anthropic Claude–4 Sonnet. These figures suggest that narratives foregrounding high power distance or strong in-group collectivism are comparatively neglected.

Between-model variability. Cluster preferences vary in magnitude, but not in direction. The coefficient of variation is smallest for **Latin Europe** ($CV = 0.09$), indicating that LLMs converge on a modest but stable share for this cluster. In contrast, **Anglo** selections display the greatest relative spread ($CV = 0.23$), hinting that “mainstream Western” advice is more LLM-specific than often assumed.

Which values do the LLMs foreground? Table 2 shows how often each system cites the nine GLOBE dimensions when justifying its selections. The dimensions are *Power Distance (PD)*, *Uncertainty Avoidance (UA)*, *Humane Orientation (HO)*, *Institutional Collectivism (IC)*, *In-Group Collectivism (IGC)*, *Assertiveness (AS)*, *Gender Egalitarianism (GE)*, *Future Orientation (FO)*, and *Performance Orientation (PO)*. Two broad patterns emerge.

1. A future-oriented, high-achievement consensus.

Across the 17 LLMs, four dimensions dominate the rationales:

Dimension	Mean Mention (%)	s.d.
In-Group Collectivism	38.2	9.5
Performance Orientation	37.1	9.4
Future Orientation	35.0	7.4
Institutional Collectivism	33.5	10.5

Table 1: The four most cited GLOBE dimensions in the rationales of 17 LLMs.

Model	PD	UA	HO	IC	IGC	AS	GE	FO	PO
jamba-1.6	21.5	31.3	28.4	43.1	18.0	1.8	2.4	25.2	25.9
claude-3.7	48.6	32.0	19.0	38.8	42.2	10.6	8.5	46.7	48.0
claude-4	41.7	43.0	43.8	52.2	23.9	8.0	13.6	48.1	50.5
command-r	19.3	23.3	32.0	29.4	37.5	7.1	6.6	22.7	20.4
deepseek	36.1	23.0	44.5	33.4	39.8	6.1	5.9	35.9	46.2
gemini-2.0	30.2	21.2	34.1	13.6	47.3	7.8	6.4	36.2	36.6
gemini-2.5	39.1	23.5	16.1	31.6	45.1	3.2	6.0	34.8	43.7
llama-3.3	35.2	19.6	37.0	24.2	24.7	5.8	5.6	44.4	38.7
llama-4	22.2	19.8	15.2	22.6	35.0	5.8	4.7	33.0	36.7
phi-4	34.1	17.2	29.1	26.6	47.4	2.7	4.3	24.3	21.6
wizardlm	27.9	23.3	25.8	30.2	35.8	6.9	6.5	30.5	50.9
mistral-nemo	26.2	20.8	24.6	19.6	42.0	7.7	4.5	34.2	31.4
gpt-4o	27.3	25.6	29.6	43.8	44.4	3.2	5.1	35.0	39.8
gpt-4.1	28.8	27.1	37.1	49.2	54.7	2.9	6.4	31.3	38.3
sonar	18.0	17.6	37.5	42.3	30.7	5.5	4.1	43.1	40.7
qwen-2.5	15.5	12.5	34.1	34.8	40.0	3.2	3.5	37.2	29.4
grok-2	20.3	17.3	31.1	34.1	41.3	2.4	3.6	33.1	32.2

Table 2: Model-wise comparison of mention rates for 9 GLOBE cultural dimensions cited in 17 LLMs rationales.

This indicates that, irrespective of the cultural cluster they ultimately pick, LLMs systematically justify their advice in terms of group solidarity, long-range thinking, and getting things done.

- Low appetite for confrontation or gender talk.** At the opposite extreme, two dimensions are almost absent: *AS*: 5.2% ($\sigma = 2.6$) and *GE*: 5.7% ($\sigma = 2.5$). No LLM mentions *AS* in more than 11% of its rationales, and the highest rate for *GE* is 13.6% (Claude-4 Sonnet). Even LLMs that frequently invoke *PD* (e.g. Claude-3 Sonnet, 48.6%) do so via “respect for hierarchy” rather than through a direct appeal to forceful negotiation.

Between-model contrasts. While the aggregate rankings are stable, individual profiles diverge: (i) **Anthropic LLMs** couple very high *PO* $\approx 50\%$ with the steepest *PD* scores, mirroring their Nordic cluster preference for meritocratic—but-still-hierarchical solutions. (ii) **Google’s Gemini 2.x** variants under-index on *IC* (13.6–31.6%) yet over-index on *HO*, suggesting a more person-centred rationale style. (iii) **OpenAI’s GPT-4.1** is an outlier for *IGC* (54.7%), framing advice in terms of familial or team loyalty more than any other system. Convergent evidence across clusters and evaluative dimensions yields a single, coherent conclusion. When confronted with ten anonymized, order-balanced alternatives, each of the 17 LLMs shows a marked preference for the normative frameworks of Northern and Western Europe, especially the Nordic and Germanic clusters, signaling an alignment imprint rather than stochastic variation. Their accompanying rationales tell the same story: although explicitly invited to invoke any of GLOBE’s nine cultural dimensions, LLMs repeatedly draw on future orientation, collective responsibility, and performance achievement, while rarely appealing to *AS* or *GE*. Taken together, the results suggest prevailing alignment pipelines embed a conciliatory, progress-centred worldview emblematic of affluent liberal democracies, functional in many settings yet

inadequate to capture humanity’s full breadth.

Are these preferences domain-specific (e.g., religion vs. business) or globally consistent? (RQ2)

To probe whether the cultural leanings uncovered in RQ1 persist across topical contexts, we split the 2,182 dilemmas into the seven CCD-Bench domains.

Dominant Nordic trend. In every domain, the **Nordic Europe** option remains the single most frequent choice. Its mean share varies from a high of 23.9% in *FAMILY* to a low of 15.6% in *LIFESTYLE*, but the ranking never drops below first place. Hence, the headline bias identified in RQ1 is not an artefact of a particular topic, such as business or science; it is a cross-domain constant of the evaluated LLMs.

Domain: Sci., Work, Educ., Wellness, Lifestyle, Family, Arts
Nordic share (%): 21.9, 20.6, 21.2, 21.0, 15.6, 23.9, 15.8
Rank 1, 1, 1, 1, 1, 1, 1

Which clusters fluctuate? While the top slot remains stable, the lower tail shifts according to context. The **least** referenced cluster varies by domain: *Eastern Europe* in *SCIENCE*, *Middle East & North Africa* in *WORK* and *EDUCATION*, *Latin America* in *WELLNESS*, *Latin Europe* in *LIFESTYLE*, and *Confucian Asia* in *ARTS*. Domain content thus nudges *which* minority perspective is discounted. **Germanic Europe** and **Sub-Saharan Africa** exhibit the widest domain variance (s.d. ≈ 3 pp), rising in *WORK* (Germanic) or *WELLNESS* (Sub-Saharan), yet receding elsewhere.

Statistical test. Collapsing all LLMs into a single 7×10 contingency matrix and applying Pearson’s χ^2 test yields $\chi^2 = 256.7$, $df = 54$, $p < .001$. However, the effect size is modest: Cramér’s $V = .12$ (95% CI [.10, .14]), labeled “small.” In other words, the domain does influence cluster selection, but the global Nordic skew significantly outweighs its magnitude.

Cultural preferences in current LLMs are *globally consistent*. Topical context shifts the edges of the distribution, deciding which clusters are least represented. Still, it does not perturb the overarching hierarchy: Nordic-style, low-power-distance, achievement-oriented advice dominates whether the user asks about science policy, workplace etiquette, family dilemmas, or the arts. The trends diagnosed in RQ1 are therefore systemic, not situational, and should be addressed at the alignment or data-curation level rather than via domain-specific patching.

Do LLMs integrate multiple cultural dimensions or default to a single axis? (RQ3)

High pluralism. Across the whole corpus ($N = 37,092$ rationales), 87.9% cite two or more dimensions, indicating that most LLMs *can* articulate multi-faceted justifications when prompted. Yet, the model spread is broad, ranging from a low of 68.6% (Meta-Llama 4 Maverick) to a peak of 98.6% (Anthropic Claude-4 Sonnet). 7 of the 17 systems exceed the 90% mark, whereas three systems remain below 80%. High pluralism is therefore common.

Pluralism \neq cultural breadth. One might expect an LLM that distributes its selections broadly (i.e., exhibits high entropy) also to draw on more dimensions, but the data suggest otherwise. The Pearson correlation between entropy and pluralism is mildly *negative* ($r = -0.22$). Cohere `command-r` exhibits the *highest* entropy (3.28) yet one of the lower pluralism scores (78.5%), while Anthropic Claude-4 shows the reverse. Integrating multiple dimensions and sampling widely across clusters are therefore distinct.

Which pairings dominate? Two results stand out: first, a single pairing explains one-fifth of all plural rationales, with *FO + PO* appearing 6,732 times (20.6% of all plural explanations), no other pairing reaches 16%, meaning LLMs overwhelmingly justify “multi-dimensional” decisions as “long-term planning to achieve goals”; second, collectivist modifiers dominate the next tier, splicing *IC*, *IGC*, or *HO* onto either *PO* or *FO*, with the top six pairings, *FO + PO* (20.6%), *HO + IGC* (15.3%), *IC + PO* (14.4%), *FO + IC* (12.9%), *IGC + PD* (12.4%), and *HO + IC* (10.7%), accounting for 86% of all plural rationales.

Dimension-level centrality. Summing across all pairings shows the frequency with which each GLOBE dimension enters multi-dimensional rationales: *PO* (16.2%), *FO* (15.1%), *IC* (14.9%), *IGC* (13.5%), *HO* (13.0%), *PD* (12.5%), *UA* (9.5%), *GE* (2.7%), and *AS* (2.6%). The first six account for over 85% of mentions, with *AS* and *GE* remaining fringe (<3% each).

Model signatures. The expanded per-model breakdown shows that: (i) **Anthropic models** over-index on the “achievement + hierarchy” tripod: *PO*, *IC*, and *PD*. (ii) **Google Gemini 2.x** variants distinguish themselves through frequent *IGC + PO* pairings, signaling a “team success” rhetoric. (iii) **Meta-Llama 4 Maverick** is the only model where *FO + PO* accounts for less than 24% of its plural rationales, yet it also exhibits the lowest pluralism rate overall.

LLMs integrate multiple cultural dimensions, but combinations are imbalanced: four mainstream values, *FO*, *PO*, and the two forms of Collectivism, dominate $\approx 90\%$ of multi-dimensional rationales, while *AS* and *GE* remain marginal. Thus, models exhibit limited pluralism, relying on a narrow subset of harmonious dimensions.

Do Models Exhibit Systematic Ordering Bias When Selecting Among Options? (RQ4)

Our generated corpus has a Latin-square balance score of $\approx 99.8\%$. Despite employing the Stratified Latin Square design to neutralize serial and residual ordering biases, these biases may persist if models exhibit systematic preferences for certain positions (e.g., favoring early or late options). In this section, we analyze the distribution of selections across the 10 positions for each cultural cluster, aggregated over all 17 models and 2,182 dilemmas (total valid selections: 36,977; expected per position under uniformity: 3,697.7). We quantify positional bias using Cramér’s V per cluster, interpreting values below 0.10 as negligible. The corresponding visualization is available in the GitHub repository.

Overall Positional Bias Levels. All 10 clusters exhibit negligible positional bias, with Cramér’s V values ranging from 0.0510 (Nordic Europe) to 0.0877 (Eastern Europe), confirming no small, medium, or large effects across the board. This low bias underscores the effectiveness of our Latin Square design in mitigating ordering artifacts. A detailed breakdown can be found on GitHub.

Mild Primacy Effect in Positions. A modest yet discernible primacy tendency surfaces: Position 1 (*P1*) is chosen disproportionately in nine of the ten cultural clusters, 15.6% for Confucian Asia and 15.5% for Eastern Europe, for example, producing an aggregate deviation of 42.2 from the uniform baseline. Nordic Europe constitutes the sole exception, exhibiting a slight preference for *P4* (12.7%) over *P1* (11.4%). Choices in the later positions (*P6–P10*) remain evenly distributed, whereas *P5* proves most balanced overall, with a deviation of only 8.6.

Cluster-Specific Variations. Clusters generally under-selected, Eastern Europe (2,086 total selections) and the Middle East (2,153), manifest the most significant positional biases ($V = 0.0877$ and $V = 0.0797$, respectively). Both clusters show elevated selection rates in early slots (Eastern Europe’s *P1* at 15.5%) and marked declines in later ones (*P8* at 6.4%). Conversely, highly favored clusters, such as Nordic Europe, register the lowest positional variance (2.34), indicating a more stable distribution across positions.

Together, these patterns confirm that the experimental protocol largely suppresses ordering effects; any residual positional bias is insufficient to compromise the cultural preference findings reported in the preceding RQs. The persistent over-selection of *P1* suggests that future benchmarks could benefit from enhanced randomization schemes.

Do LLMs Cluster by Cultural Behavior, and Does This Pattern Correlate with Their Geographic Provenance? (RQ5)

To assess similarities in cultural selection patterns and their relation to developers’ locations, we compute symmetrized Kullback-Leibler (KL) divergence between each pair of models’ cluster distributions, yielding a 17×17 distance matrix fed into average-linkage hierarchical clustering (See Figure 2). Developer geographies are from public sources.

Model Clustering Patterns. The dendrogram reveals tight intra-family groupings at low divergence levels (e.g., Meta’s Llama-3.3 and Llama-4 at 0.0033, Anthropic’s Claude-3.7 and Claude-4 at 0.0398, Google’s Gemini-2.x at 0.0103, OpenAI’s GPT-4o at 0.0200, Microsoft’s Phi-4 and GPT-4.1 at 0.0032 via linkage). Broader clusters emerge around 0.01–0.05: one subgroup includes Google, OpenAI, xAI, Cohere, and Qwen; another features Anthropic, Meta, Perplexity, Mistral, and DeepSeek.

Correlation with Geographic Provenance. Clustering shows weak geographic ties: US models span both sub-clusters, with corporate lineage (e.g., Meta/Anthropic vs. Google/OpenAI) dominating over location. Non-US models fragment similarly (Chinese: DeepSeek in subcluster




Figure 2: Heatmap of KL-divergence between LLMs' cultural cluster selections (left), with dendrogram (right).

2, Qwen in 1; European-adjacent: Mistral in 2, Cohere in 1; AI21 with Google/Cohere). This implies that alignment strategies and shared corpora influence cultural behavior more than national origin, though Western-centric data limits distinctions.

Discussion

Our findings reveal a clear cultural bias in contemporary LLMs: when presented with plural legitimate options in a decision-making task, models tend to gravitate toward Nordic/Germanic clusters. The associated rationals are strikingly uniform: models rely on *Future Orientation* and *Performance Orientation* with regularity, while rarely grounding decisions in *Assertiveness* or *Gender Egalitarianism* (each < 3%). This does not merely imply that models have bias; rather, the very *style* of advice being normalised in highly used models embodies a conciliatory, progress-centred, achievement-first worldview, which is systematically misaligned with contexts that legitimate hierarchy or where an explicit appeal to gender fairness is the appropriate move. Therefore, occasionally, a default to consensus-seeking performance talk can result in advice that comes across as naive or disrespectful. Models should instead surface plural options, making trade-offs explicit and labeling the value frames that motivate each option, allowing users to decide which options are valid given their local norms.

A notable limitation is our exclusive use of English prompts, which may bias outputs toward Anglo-American values and obscure regional nuances (Zhong, Yun, and Sun 2024); future work should utilize multilingual prompts to distinguish between language and cultural effects. If LLMs are to mediate everyday social decision-making, alignment must treat *conflicts among legitimate values* as a design target, so that a model helping someone write to their supervisor or address a family dispute can offer guidance that is not simply *good*, but *good for their values*.

Conclusion

In this work, we introduce *CCD-Bench*, a novel benchmark specifically designed to assess how contemporary LLMs navigate conflicts among legitimate cultural values. Evaluating 17 LLMs on 2,182 dilemmas, we observed a pronounced cultural trend toward Nordic and Germanic clusters. Despite high rates of multi-dimensional rationales, models relied on a narrow set of GLOBE dimensions, while rarely invoking *Assertiveness* or *Gender Egalitarianism* (<3% each), suggesting an aversion to confrontational or gender-focused reasoning and indicating limited true pluralism in reasoning. Our results underscore the need for more nuanced approaches to cultural dynamics in LLMs if they are to operate credibly within an increasingly pluralistic world.

Acknowledgments

This work is supported in part by the NYUAD Center for Interdisciplinary Data Science & AI, funded by Tamkeen under the NYUAD Research Institute Award CG016.

References

- Agresti, A. 2019. *An introduction to categorical data analysis*. Wiley. ISBN 978-1-119-40526-9.
- AlKhamissi, B.; ElNokrashy, M.; Alkhamissi, M.; and Diab, M. 2024. Investigating Cultural Alignment of Large Language Models. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 12404–12422. Bangkok, Thailand: Association for Computational Linguistics.
- Askill, A.; Bai, Y.; Chen, A.; Drain, D.; Ganguli, D.; Henighan, T.; Jones, A.; Joseph, N.; Mann, B.; DasSarma, N.; Elhage, N.; Hatfield-Dodds, Z.; Hernandez, D.; Kernion, J.; Ndousse, K.; Olsson, C.; Amodei, D.; Brown, T.; Clark, J.; McCandlish, S.; Olah, C.; and Kaplan, J. 2021. A General Language Assistant as a Laboratory for Alignment. ArXiv:2112.00861 [cs].
- Bai, Y.; Jones, A.; Ndousse, K.; Askill, A.; Chen, A.; DasSarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; Joseph, N.; Kadavath, S.; Kernion, J.; Conerly, T.; El-Showk, S.; Elhage, N.; Hatfield-Dodds, Z.; Hernandez, D.; Hume, T.; Johnston, S.; Kravec, S.; Lovitt, L.; Nanda, N.; Olsson, C.; Amodei, D.; Brown, T.; Clark, J.; McCandlish, S.; Olah, C.; Mann, B.; and Kaplan, J. 2022. Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback. ArXiv:2204.05862 [cs].
- Balepur, N.; Rudinger, R.; and Boyd-Graber, J. L. 2025. Which of These Best Describes Multiple Choice Evaluation with LLMs? A) Forced B) Flawed C) Fixable D) All of the Above. ArXiv:2502.14127 [cs].
- Barbara Herrnstein Smith. 1991. *Contingencies of value : alternative perspectives for critical theory*. Harvard University Press. ISBN 978-0-674-16786-5.
- Bradley, J. V. 1958. Complete Counterbalancing of Immediate Sequential Effects in a Latin Square Design. *Journal of the American Statistical Association*, 53(282): 525–528.
- Cann, T. J. B.; Dennes, B.; Coan, T.; O’Neill, S.; and Williams, H. T. P. 2025. Using semantic similarity to measure the echo of strategic communications. *EPJ Data Science*, 14(1): 1–19. Number: 1 Publisher: SpringerOpen.
- Chiu, Y. Y.; Jiang, L.; Lin, B. Y.; Park, C. Y.; Li, S. S.; Ravi, S.; Bhatia, M.; Antoniak, M.; Tsvetkov, Y.; Shwartz, V.; and Choi, Y. 2024. CulturalBench: a Robust, Diverse and Challenging Benchmark on Measuring the (Lack of) Cultural Knowledge of LLMs. ArXiv:2410.02677 [cs] version: 1.
- Dorfman, R. 1979. A Formula for the Gini Coefficient. *The Review of Economics and Statistics*, 61(1): 146.
- Efron, B.; and Tibshirani, R. 1998. *An introduction to the bootstrap*. Boca Raton, Fla Chapman & Hall/Crc. ISBN 978-0-412-04231-7.
- Elgarf, M.; Salam, H.; and Peters, C. 2024. Fostering children’s creativity through LLM-driven storytelling with a social robot. *Frontiers in Robotics and AI*, 11: 1457429.
- Harald Cramér. 1999. *Mathematical methods of statistics*. Princeton University Press. ISBN 978-0-691-00547-8.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2021. Measuring Massive Multitask Language Understanding. ArXiv:2009.03300 [cs].
- House, R. J.; Hanges, P. J.; Javidan, M.; Dorfman, P. W.; and Gupta, V. 2004. *Culture, leadership, and organizations: The GLOBE study of 62 societies*. Sage Publications, 1 edition.
- Karinshak, E.; Hu, A.; Kong, K.; Rao, V.; Wang, J.; Wang, J.; and Zeng, Y. 2024. LLM-GLOBE: A Benchmark Evaluating the Cultural Values Embedded in LLM Output. ArXiv:2411.06032 [cs].
- Koçak, Ö.; Puranam, P.; and Yeğın, A. 2024. LLMs as mediators: Can they diagnose conflicts accurately? ArXiv:2412.14675 [cs].
- Kullback, S. 1997. *Information theory and statistics*. Dover Publications. ISBN 978-0-486-69684-3.
- Lalwani, H.; Saleh, M.; and Salam, H. 2025. A study companion for productivity: exploring the role of a social robot for college students with ADHD. In *2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 1438–1442. IEEE.
- Lu, J. G.; Song, L. L.; and Zhang, L. D. 2025. Cultural tendencies in generative AI. *Nature Human Behaviour*, 1–10. Publisher: Nature Publishing Group.
- Myung, J.; Lee, N.; Zhou, Y.; Jin, J.; Putri, R. A.; Antypas, D.; Borkakoty, H.; Kim, E.; Perez-Almendros, C.; Ayele, A. A.; Gutiérrez-Basulto, V.; Ibáñez-García, Y.; Lee, H.; Muhammad, S. H.; Park, K.; Rzayev, A. S.; White, N.; Yimam, S. M.; Pilehvar, M. T.; Ousidhoum, N.; Camacho-Collados, J.; and Oh, A. 2024. BLEnD: A Benchmark for LLMs on Everyday Knowledge in Diverse Cultures and Languages. *Advances in Neural Information Processing Systems*, 37: 78104–78146.
- Naous, T.; Ryan, M. J.; Ritter, A.; and Xu, W. 2024. Having Beer after Prayer? Measuring Cultural Bias in Large Language Models. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 16366–16393. Bangkok, Thailand: Association for Computational Linguistics.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C. L.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS ’22*, 27730–27744. Red Hook, NY, USA: Curran Associates Inc. ISBN 978-1-7138-7108-8.
- Pezeshkpour, P.; and Hruschka, E. 2024. Large Language Models Sensitivity to The Order of Options in Multiple-Choice Questions. In Duh, K.; Gomez, H.; and Bethard,

- S., eds., *Findings of the Association for Computational Linguistics: NAACL 2024*, 2006–2017. Mexico City, Mexico: Association for Computational Linguistics.
- Pires, A. S.; Samson, L.; Ghebreab, S.; and Santos, F. P. 2025. How large language models judge and influence human cooperation. ArXiv:2507.00088 [physics].
- Shannon, C. E.; and Weaver, W. 1949. The Mathematical Theory of Communication. *The Mathematical Gazette*, 34(310): 312.
- Sneath, P. H.; and Sokal, R. R. 1973. *Numerical Taxonomy: The Principles and Practice of Numerical Classification*. W. H. Freeman.
- Tan, J.; Westermann, H.; Pottanigari, N. R.; Šavelka, J.; Meeüs, S.; Godet, M.; and Benyekhlef, K. 2024. Robots in the Middle: Evaluating LLMs in Dispute Resolution. ArXiv:2410.07053 [cs].
- Wang, Y.; Zhu, Y.; Kong, C.; Wei, S.; Yi, X.; Xie, X.; and Sang, J. 2024. CDEval: A Benchmark for Measuring the Cultural Dimensions of Large Language Models. In Prabhakaran, V.; Dev, S.; Benotti, L.; Hershcovich, D.; Cabello, L.; Cao, Y.; Adebara, I.; and Zhou, L., eds., *Proceedings of the 2nd Workshop on Cross-Cultural Considerations in NLP*, 1–16. Bangkok, Thailand: Association for Computational Linguistics.
- Xu, S.; Leng, Y.; Yu, L.; and Xiong, D. 2025. Self-Pluralising Culture Alignment for Large Language Models. In Chiruzzo, L.; Ritter, A.; and Wang, L., eds., *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 6859–6877. Albuquerque, New Mexico: Association for Computational Linguistics. ISBN 979-8-89176-189-6.
- Zhao, W.; Mondal, D.; Tandon, N.; Dillion, D.; Gray, K.; and Gu, Y. 2024. WorldValuesBench: A Large-Scale Benchmark Dataset for Multi-Cultural Value Awareness of Language Models. ArXiv:2404.16308 [cs].
- Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E. P.; Zhang, H.; Gonzalez, J. E.; and Stoica, I. 2023. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. ArXiv:2306.05685 [cs].
- Zhong, Q.; Yun, Y.; and Sun, A. 2024. Cultural Value Differences of LLMs: Prompt, Language, and Model Size. ArXiv:2407.16891 [cs] version: 1.