



I AM ALIGNED, BUT WITH WHOM? MENA VALUES BENCHMARK FOR EVALUATING CULTURAL ALIGNMENT AND MULTILINGUAL BIAS IN LLMs

Anonymous authors

Paper under double-blind review

ABSTRACT

We introduce MENAValues, a novel benchmark designed to evaluate the cultural alignment and multilingual biases of large language models (LLMs) with respect to the beliefs and values of the Middle East and North Africa (MENA) region, an underrepresented area in current AI evaluation efforts. Drawing from large-scale, authoritative human surveys, we curate a structured dataset that captures the sociocultural landscape of MENA with population-level response distributions from 16 countries. To probe LLM behavior, we evaluate diverse models across multiple conditions formed by crossing three perspective framings (neutral, personalized, and third-person/cultural observer) with two language modes (English and localized native languages: Arabic, Persian, Turkish). Our analysis reveals three critical phenomena: "Cross-Lingual Value Shifts" where identical questions yield drastically different responses based on language, "Reasoning-Induced Degradation" where prompting models to explain their reasoning worsens cultural alignment, and "Logit Leakage" where models refuse sensitive questions while internal probabilities reveal strong hidden preferences. We further demonstrate that models collapse into simplistic linguistic categories when operating in native languages, treating diverse nations as monolithic entities. MENAValues offers a scalable framework for diagnosing cultural misalignment, providing both empirical insights and methodological tools for developing more culturally inclusive AI.

1 INTRODUCTION

Large language models (LLMs) have achieved remarkable capabilities, yet they often struggle to align with the nuanced cultural norms and values of diverse global communities Shen et al. (2024b); Li et al. (2024b). This misalignment stems largely from training datasets that predominantly reflect Western and English-speaking perspectives, resulting in a constrained understanding of culturally grounded knowledge Durmus et al. (2024b); Naous et al. (2024c). The Middle East and North Africa (MENA) region, comprising over 500 million people across 16+ countries, epitomizes this challenge of cultural misrepresentation. While recent research has made strides in assessing cultural values and biases in LLMs Mitchell et al. (2025); Tao et al. (2024b); Kharchenko et al. (2024), empirical studies systematically examining how these models represent MENA populations' values remain scarce. The implications of this gap are profound. As LLMs are increasingly deployed globally, their ability to accurately reflect diverse cultural perspectives directly impacts user experiences, trust, and the potential for intercultural understanding. Our research addresses critical limitations in current LLM cultural alignment approaches through the MENAValues Benchmark, a comprehensive evaluation framework with three key innovations:

- A novel dataset of 864 questions derived from population-scale survey data, covering diverse aspects of MENA cultural values and beliefs.
- A robust evaluation methodology that examines LLM responses across multiple prompting conditions, languages, and models.
- An analytical framework that explores token-level probabilities to reveal hidden biases and internal contradictions, uncovering novel phenomena such as Logit Leakage and Reasoning-Induced Degradation.




Figure 1: **Systematic Value Inconsistency in LLMs: A Multi-Dimensional Analysis of Alignment Failures.** This figure reveals how LLMs exhibit inconsistencies when responding to value-based questions, demonstrating three critical dimensions of misalignment. **Cross-Lingual Value Shift** shows how identical questions yield contradictory responses across languages (Arabic vs. English), suggesting cultural bias encoding rather than consistent moral reasoning. **Prompt-Sensitive Misalignment** reveals how different framings (Neutral, Persona, Observer) elicit conflicting stances, indicating unstable value representation. **Logit Leakage** exposes how surface-level safety responses can mask concerning internal preferences, with high-confidence hidden biases contradicting stated positions.

The benchmark interrogates several crucial questions about LLMs' cultural alignment: How accurately do current models reflect the documented values of MENA populations? How do language and identity framing alter model responses? Do models exhibit contradictions between their expressed outputs and internal probability distributions? By systematically investigating these questions, we provide unprecedented insights into AI alignment challenges in multilingual and culturally diverse contexts. Our approach transforms data from the World Values Survey Wave 7 Haerpfer et al. (2022) and the 2022 Arab Opinion Index Arab Center for Research and Policy Studies (2022) into multiple-choice questions spanning governance, economics, cultural identity, and individual wellbeing. We evaluate LLMs across a matrix of conditions, varying perspective (neutral, personalized, and observer) and language (English versus native languages including Arabic, Persian, and Turkish). As demonstrated in Figure 1, our research reveals persistent inconsistencies in model behavior, highlighting three critical misalignment behaviors. These findings raise significant concerns about LLM reliability in multilingual and culturally sensitive contexts. Our work addresses these gaps by advancing culturally inclusive AI alignment research and providing a methodological template for evaluating cultural alignment across underrepresented regions, building upon prior research AlKhamissi et al. (2024b); Durmus et al. (2024b).

2 RELATED WORK

Recent work on cultural alignment in LLMs has revealed significant gaps in representing diverse cultural perspectives Li et al. (2024a); Kirk et al. (2024); Durmus et al. (2024a); AlKhamissi et al. (2024a); Ryan et al. (2024); Gabriel & Ghazavi (2021); Wang et al. (2024b); Adilazuarda et al. (2024). Cross-cultural NLP research has identified persistent challenges in addressing cultural nuances Hershovich et al. (2022b;a), with studies documenting biases against Muslim and Arab communities Naous et al. (2024b;a); Abid et al. (2021). This has motivated the development of benchmarks including StereoSet Nadeem et al. (2021), StereoKG Deshpande et al. (2022), SEEGULL Jha et al. (2023), and CultureBank Shi et al. (2024), alongside frameworks for auditing cultural biases Tao et al. (2024a); Gupta et al. (2024); Sheng et al. (2021). Several projects have addressed specific regional needs, including Arabic localization Huang et al. (2024) and frameworks for evaluating regional

108 cultural reasoning Cao et al. (2023); Fung et al. (2024); Wang et al. (2024a). Concurrently, research
 109 has examined cross-cultural differences in values Arora et al. (2023) and approaches to align AI
 110 with diverse human values Hendrycks et al. (2023), while empirical studies have assessed alignment
 111 between language models and various cultural contexts Cao et al. (2023); Wang et al. (2024b); Arora
 112 et al. (2023).

113 The multilingual capabilities of LLMs profoundly impact cultural representation, with studies re-
 114 vealing disparities in performance across languages Etxaniz et al. (2024) and inconsistencies in
 115 factual knowledge Qi et al. (2023) and safety behaviors Shen et al. (2024a) across linguistic con-
 116 texts. Solutions to these challenges include modular transformer architectures Pfeiffer et al. (2022),
 117 language-neutral sub-networks Foroutan et al. (2022), and improved cross-lingual consistency eval-
 118 uation frameworks Qi et al. (2023); Wang et al. (2024a). The effect of anthropomorphism and
 119 persona-based evaluation on model outputs has gained increasing attention Deshpande et al. (2023);
 120 Joshi et al. (2024); Kirk et al. (2023); Jang et al. (2023); Cheng et al. (2024; 2023), revealing how
 121 identity cues and framing significantly impact model responses when addressing culturally sensitive
 122 topics. Work on understanding and mitigating social biases in language models Liang et al. (2021)
 123 has shown the importance of examining both explicit outputs and underlying patterns Deshpande
 124 et al. (2023); Joshi et al. (2024); Cheng et al. (2023).

125 3 MENAVALUES BENCHMARK

126 The MENAValues benchmark is designed to evaluate the cultural alignment and multilingual behavior
 127 of LLMs in the context of the MENA region. It provides a structured and diverse set of multiple-
 128 choice questions that reflect real-world human values with a spectrum of choices to select the stance.
 129 Organized into topical categories and covering 16 MENA countries, the benchmark preserves human
 130 response distributions and enables fine-grained analysis of model responses under varying linguistic
 131 and perspective-based prompting conditions.

132 3.1 DATA SOURCES

133 The MENAValues benchmark is constructed from two large-scale, high-quality survey datasets: the
 134 *World Values Survey Wave 7 (WVS-7)* and the *Arab Opinion Index 2022 (AOI-2022)*. Both instruments
 135 offer nationally representative samples and cover a wide range of sociopolitical, economic, and
 136 cultural attitudes.

137 **WVS-7** is a global survey effort conducted between 2017 and 2022, with responses collected from
 138 over 77 countries using standardized questionnaires. It includes 291 questions, covering both universal
 139 themes and 32 region-specific items for MENA countries. From WVS-7, we include data from nine
 140 MENA countries: Egypt, Iran, Iraq, Jordan, Lebanon, Libya, Morocco, Turkey, and Tunisia. The
 141 questions span 14 thematic domains, including social values, well-being, political participation, trust,
 142 corruption, and migration.

143 **AOI-2022** is a regionally focused survey conducted by the Arab Center for Research and Policy
 144 Studies, encompassing 573 questions from 14 Arab countries. It captures public opinion on a
 145 wide range of issues such as governance, civil liberties, state performance, regional politics, and
 146 personal dignity. Countries represented include Algeria, Egypt, Iraq, Jordan, Kuwait, Lebanon, Libya,
 147 Mauritania, Morocco, Palestine, Qatar, Saudi Arabia, Sudan, and Tunisia.

148 Together, these sources provide a rich, empirically grounded foundation of 864 total questions across
 149 16 MENA countries. We selected questions based on their thematic relevance, clarity, and cross-
 150 national applicability. To ensure our ground-truth data accurately reflects national demographics, we
 151 apply post-stratification weights provided by the surveys.

152 3.2 REGIONAL COHERENCE AND DIVERSITY ANALYSIS

153 To examine MENA as a coherent yet internally diverse cultural region, we conducted compre-
 154 hensive analysis of variance patterns across countries and thematic categories (detailed analysis
 155 in Appendix C). Critically, Jensen-Shannon Divergence analysis demonstrates exceptionally high
 156 distributional similarity scores (>0.95) between country pairs, indicating a shared "grammar" of
 157 opinion expression across the region despite substantive disagreements on specific issues. Thematic




Figure 2: **Core Dimensions of the MENAVALUES Dataset.** The dataset is structured around four major pillars: (1) *Social & Cultural Identity*, (2) *Economic Dimensions*, (3) *Governance & Political Systems*, and (4) *Individual Wellbeing & Development*. Each category is illustrated with survey questions and average responses from representative MENA countries. Note: The countries shown are illustrative examples; all benchmark questions are posed to LLMs for all 16 countries in our dataset, regardless of original human data availability.

variance analysis reveals important patterns: Individual Wellbeing & Development shows the highest regional consensus (variance: 0.1297 in WVS, 0.3242 in AOI), while Economic Dimensions (WVS) and Social & Cultural Identity (AOI) exhibit the greatest divergence. These patterns shift depending on whether non-Arab countries are included, with primary fault lines moving from economic to identity-based divisions within the exclusively Arab sample.

3.3 BENCHMARK CATEGORIES

To structure the benchmark, we organize questions into four broad topical categories that reflect central dimensions of public values in the MENA region: *Governance & Political Systems*, *Economic Dimensions*, *Social & Cultural Identity*, and *Individual Wellbeing & Development*. This taxonomy enables both high-level and fine-grained evaluation of model behavior across distinct sociopolitical and ethical domains. Each category aggregates multiple subtopics derived from the original survey structure. From WVS-7, 291 questions are distributed across these categories as follows: 36.8% in governance, 47.8% in social and cultural identity, 9.6% in individual wellbeing, and 5.8% in economic dimensions. AOI-2022 provides a complementary distribution, with 82.2% of questions in governance, 12.7% in wellbeing, and 5.1% in cultural identity. Figure 2 illustrates the benchmark’s core structure, presenting the four principal dimensions along with representative subcategories and example questions. It also displays average human responses and country-wise distributions, as well as corresponding outputs and logit probabilities from LLMs.

3.4 QUESTION FORMULATION

To stay true to the original surveys, we maintain the multiple-choice and Likert-scale formats to transform survey items into a machine-evaluable benchmark. We also capture full response distributions and token-level probabilities to enhance our analysis beyond the constraints typically

216 associated with discrete choice formats Li et al. (2024c); Zheng et al. (2024); Balepur et al. (2024).
 217 To derive a final human ground-truth estimate from the distribution of thousands of responses, we use
 218 the majority response for categorical questions and compute the weighted mean response for scalar
 219 (spectrum-based) items. These aggregated responses serve as representative benchmarks against
 220 which LLM predictions and responses from other demographic groups can be compared.

221 Crucially, we retain the full human response distributions to support deeper analyses, such as
 222 comparing the model’s output probability distribution to the human distribution using KL divergence,
 223 as well as downstream tasks such as logit-probability analysis and refusal behavior modeling. Each
 224 item is thus encoded with both its representative answer and its underlying distribution, enabling
 225 evaluation of LLM alignment with public opinion across multiple analytical dimensions.

227 4 EVALUATION FRAMEWORK

230 We present a multidimensional evaluation framework for assessing how LLMs represent MENA
 231 cultural values across varying perspectives, languages, and reasoning conditions. Our methodology
 232 thoroughly examines framing effects through three distinct prompting styles, two language conditions,
 233 and two reasoning conditions to uncover potential alignment failures and cultural biases. This creates
 234 a comprehensive $3 \times 2 \times 2$ evaluation matrix that enables fine-grained analysis of how different factors
 235 interact to shape cultural representation in LLMs.

237 4.1 PERSPECTIVE AXIS (FRAMING STYLES)

239 We design our evaluation framework to systematically probe how framing affects LLMs’ representa-
 240 tion of MENA values, grounded in the psychological paradigm of self-report versus observer ratings,
 241 where self-reports may reflect idealized views and observer ratings may incorporate stereotypes. Our
 242 methodology incorporates three distinct prompting perspectives, each designed to elicit different
 243 aspects of model behavior:

244 **Neutral Framing:** We query the LLM directly without imposing identity constraints or cultural
 245 framing. Prompts follow the format: “[Question about value/belief]”. This baseline captures the
 246 model’s default positioning on MENA-related values and serves as our control condition.

247 **Persona-Based Framing:** We instruct the model to embody a specific national identity through
 248 prompts structured as: “*Imagine you are an average [nationality]. [Question about value/belief]*”.
 249 This anthropomorphized condition tests how explicit identity cues activate the model’s internalized
 250 cultural schemas and alter response patterns, revealing the model’s internal representations of specific
 251 MENA cultural identities.

252 **Cultural Observer Framing:** We position the model as an external analyst through prompts like:
 253 “*How would an average [nationality] respond to [question about value/belief]*?” This framing elicits
 254 the model’s sociological generalizations and understanding of MENA populations without requiring
 255 direct role-playing, potentially activating stereotype-based reasoning.

256 These framing distinctions are particularly important for culturally-nuanced topics where models
 257 might (1) default to Western-centric values in neutral framing, (2) attempt to perform culturally-
 258 specific values when role-playing, or (3) employ stereotypes when describing cultures from an
 259 observer perspective. By comparing responses across these three frames, we can identify inconsisten-
 260 cies that reveal potential alignment failures and cultural biases.

262 4.2 LANGUAGE AXIS

264 We evaluate LLMs across two linguistic conditions to assess how language choice influences cultural
 265 representation:

267 **English Prompting:** All benchmark questions are presented in English, regardless of the cultural
 268 context being probed. This condition serves as the standard evaluation approach in most LLM
 269 benchmarks and reveals how models represent MENA values when operating in the dominant
 language of AI development.

270 **Native Language Prompting:** We translate all prompts into the primary language of each respective
 271 MENA region: Arabic (for most Arab countries), Persian (for Iran), and Turkish (for Turkey). To
 272 ensure quality and cultural nuance, translations were validated by native, in-house human annotators.
 273

274 Our cross-lingual evaluation highlights three key points. First, it isolates language as a causal factor
 275 in cultural representation, testing whether models preserve consistent value frameworks or display
 276 linguistic determinism. Second, it reflects the practical reality that LLMs in MENA contexts often
 277 operate in local languages, making such evaluation critical for real-world alignment. Third, it reveals
 278 biases in multilingual training data, as divergences between English and native responses may signal
 279 uneven representation of cultural values.

280 4.3 REASONING CONDITIONS

281 To investigate how explicit reasoning processes affect cultural alignment, we evaluate models under
 282 two distinct cognitive conditions that probe different aspects of model behavior:

283 **Zero-Shot Response:** Models provide direct answers to value-based questions without additional
 284 reasoning prompts. This condition captures the model’s immediate, intuitive responses.

285 **With-Reasoning Response:** Models are instructed to provide brief reasoning before answering, this
 286 condition tests whether deliberative processes improve or degrade cultural representation.

287 The inclusion of reasoning conditions addresses a fundamental assumption in AI alignment research:
 288 that encouraging models to “think through” their responses leads to better outcomes. Our framework
 289 enables comprehensive comparison of immediate versus deliberative responses across cultural
 290 contexts, revealing whether reasoning processes activate beneficial cultural knowledge or problematic
 291 biases and stereotypes.

292 4.4 EVALUATION MODELS

293 We evaluate seven diverse LLMs: Llama-3.1-8B-Instruct Grattafiori et al. (2024) and Mistral-7B-
 294 Instruct-v0.3 Jiang et al. (2023) as general-purpose foundation models; AYA (aya-expanse-8b) Dang
 295 et al. (2024) as a multilingual-focused model with significant Arabic training data; Fanar-1-9B-Instruct
 296 Team et al. (2025) and ALLAM-Thinking Research (2025) as Arabic-centric regional specialists; and
 297 GPT-4o-mini OpenAI (2025) and Gemini 2.5 Flash Lite Google Cloud (2025) as frontier proprietary
 298 models. This selection ensures robust findings across varying parameter counts, training approaches,
 299 and cultural specializations, with most models being open-source to enable analysis of internal token
 300 probabilities.

301 5 METHODS

302 This section outlines our methodological framework for evaluating the cultural alignment of LLMs
 303 with respect to the MENA region’s values and beliefs. We present the evaluation metrics employed to
 304 quantify alignment across different dimensions, followed by our analytical approaches for examining
 305 model behavior.

306 5.1 EVALUATION METRICS

307 To comprehensively assess how accurately and consistently LLMs represent MENA values, we
 308 developed a suite of quantitative metrics that capture different aspects of alignment and model
 309 behavior. Let $Q = \{q_1, q_2, \dots, q_n\}$ denote our set of benchmark questions, $M = \{m_1, m_2, \dots, m_k\}$
 310 represent our evaluated models, and $C = \{c_1, c_2, \dots, c_l\}$ denote the set of MENA countries. For each
 311 question $q \in Q$, let $O_q = \{o_1, o_2, \dots, o_{|O_q|}\}$ represent the set of possible response options.

312 5.1.1 NORMALIZED VALUE ALIGNMENT SCORE (NVAS)

313 The Normalized Value Alignment Score measures **cultural authenticity** by assessing the degree
 314 to which model predictions align with ground truth human values from survey data. For a model
 315 $m \in M$, country $c \in C$, and question $q \in Q$, we define:

324
325
326
327

$$\text{NVA}S_{m,c} = \frac{1}{|Q|} \sum_{q \in Q} \left(1 - \frac{|v_{m,q} - v_{c,q}|}{v_{\max} - v_{\min}} \right) \times 100\% \quad (1)$$

328 where $v_{m,q}$ represents the model's predicted value for question q , $v_{c,q}$ represents the human ground
329 truth value for country c and question q , and v_{\max} , v_{\min} denote the maximum and minimum possible
330 values across all questions. This metric scales the deviation between model and human values to
331 percentages, with 100% indicating perfect alignment and 0% maximum misalignment.

332 5.1.2 CONSISTENCY METRICS FRAMEWORK

333 We employ a unified mathematical framework for measuring different aspects of model consistency.
334 Let $\mathcal{D}(v_1, v_2)$ represent the normalized distance function. Then our consistency metrics are defined
335 as:

336 **Framing Consistency Score (FCS)** Tests **cognitive coherence** by quantifying consistency across
337 different prompting perspectives:

$$\text{FCS}_{m,c} = \frac{1}{|Q|} \sum_{q \in Q} (1 - \mathcal{D}(v_{m,q}^{\text{persona}}, v_{m,q}^{\text{observer}})) \times 100\% \quad (2)$$

338 **Cross-Lingual Consistency Score (CLCS)** Evaluates **cultural universalism** by measuring consistency
339 between representations in different languages:

$$\text{CLCS}_{m,c} = \frac{1}{|Q|} \sum_{q \in Q} (1 - \mathcal{D}(v_{m,q}^{\text{English}}, v_{m,q}^{\text{Native}})) \times 100\% \quad (3)$$

340 **Self-Persona Deviation (SPD)** Captures **anthropomorphic responsiveness** by quantifying response
341 changes under persona assignment:

$$\text{SPD}_{m,c} = \frac{1}{|Q|} \sum_{q \in Q} (1 - \mathcal{D}(v_{m,q}^{\text{neutral}}, v_{m,q}^{\text{persona}})) \times 100\% \quad (4)$$

342 All reported results include 95% bootstrap confidence intervals computed over $B = 1,000$ resamples
343 to quantify uncertainty.

344 5.2 ANALYSIS APPROACHES

345 5.2.1 TOKEN PROBABILITY ANALYSIS

346 To examine model behavior beyond surface responses, we analyze the token-level probabilities
347 assigned to answer options. We extract the normalized log-probabilities for each option. This enables
348 detection of **logit leakage**, where a model with strong internal preferences refuses to provide explicit
349 answers. We define a "strong internal conviction" as any option with a normalized log-probability
350 exceeding 75% of the maximum. Alignment between model probability distributions and human
351 responses is measured using Kullback-Leibler divergence.

352 5.2.2 ABSTENTION AND REFUSAL ANALYSIS

353 We systematically track instances where models decline to provide direct answers, categorizing
354 these as "refusals." By analyzing abstention rates across different conditions, we identify patterns
355 in when models exercise caution regarding cultural judgments. This analysis reveals how different
356 prompting conditions, languages, and reasoning requirements affect models' willingness to engage
357 with culturally sensitive topics.

358 5.2.3 STRUCTURAL REPRESENTATION ANALYSIS

359 To examine the underlying organization of cultural representations beyond surface metrics, we
360 conducted Principal Component Analysis on model responses across countries and conditions. This

378 dimensional reduction technique reveals how models cluster and differentiate between cultural
 379 contexts, exposing patterns in representational structure that may not be apparent in aggregate
 380 alignment scores.

382 6 RESULTS

384 Our evaluation, spanning an immense dataset of over 820,000 data points, uncovers systematic
 385 and pervasive failures in the cultural alignment of LLMs. The findings presented here challenge
 386 fundamental assumptions about their global applicability, revealing complex patterns of cultural
 387 representation that simple performance metrics fail to capture. The complete results are summarized
 388 in Table 1 and visualized in Appendix Figure 7.

390 Table 1: Overall Evaluation Metrics Across Models and Reasoning Conditions. All values are
 391 percentages, with 95% confidence intervals in brackets. KLD is Kullback–Leibler Divergence.

Model	Reasoning	CLCS	FCS	KLD	NVAS	SPD
AYA	Zero-Shot	80.49 [79.24, 81.87]	79.18 [77.18, 80.94]	1.63 [1.60, 1.66]	70.12 [69.78, 70.45]	79.59 [77.89, 81.13]
	With-Reasoning	79.05 [78.33, 79.83] ↓	80.91 [80.16, 81.66] ↑	1.59 [1.57, 1.61] ↑	69.92 [69.68, 70.18] ↓	79.01 [78.16, 79.85] ↓
Mistral	Zero-Shot	66.54 [65.56, 67.39]	88.51 [87.46, 89.53]	2.98 [2.95, 3.02]	69.15 [68.87, 69.47]	87.21 [86.47, 87.98]
	With-Reasoning	65.44 [64.69, 66.16] ↓	83.93 [83.27, 84.59] ↓	3.56 [3.54, 3.59] ↓	65.63 [65.38, 65.89] ↓	83.04 [82.46, 83.66] ↓
Llama-3.1	Zero-Shot	79.30 [78.70, 79.88]	85.83 [85.34, 86.37]	1.31 [1.30, 1.32]	75.75 [75.55, 75.96]	83.61 [82.96, 84.26]
	With-Reasoning	70.96 [70.23, 71.61] ↓	76.55 [75.98, 77.15] ↓	1.07 [1.06, 1.08] ↑	68.79 [68.55, 69.04] ↓	74.94 [74.23, 75.64] ↓
GPT-4o-mini	Zero-Shot	89.47 [89.07, 89.89]	90.52 [90.12, 90.93]	N/A	75.34 [75.13, 75.54]	80.65 [80.08, 81.22]
	With-Reasoning	89.93 [89.55, 90.32] ↑	91.61 [91.23, 91.98] ↑	N/A	75.24 [75.05, 75.43] ↓	79.72 [79.15, 80.36] ↓
ALLAM	Zero-Shot	80.19 [77.58, 82.08]	88.85 [88.35, 89.30]	1.35 [1.34, 1.37]	70.56 [70.02, 71.04]	85.34 [84.81, 85.90]
	With-Reasoning	81.98 [81.39, 82.60] ↑	88.18 [87.71, 88.61] ↓	1.18 [1.16, 1.20] ↑	71.09 [70.85, 71.34] ↑	74.38 [73.39, 75.30] ↓
Fanar	Zero-Shot	83.10 [82.56, 83.66]	91.38 [90.93, 91.85]	2.97 [2.94, 2.99]	72.95 [72.72, 73.17]	83.51 [82.90, 84.10]
	With-Reasoning	67.09 [66.31, 67.79] ↓	71.10 [70.33, 71.89] ↓	2.84 [2.81, 2.86] ↑	66.83 [66.61, 67.07] ↓	77.38 [76.68, 78.08] ↓
Gemini	Zero-Shot	88.38 [87.91, 88.79]	89.18 [87.81, 90.20]	N/A	74.74 [74.42, 75.03]	76.80 [76.19, 77.45]
	With-Reasoning	86.98 [86.49, 87.41] ↓	85.49 [85.03, 85.98] ↓	N/A	72.32 [72.10, 72.54] ↓	73.62 [72.97, 74.31] ↓

404 KLD is not available (N/A) for closed-source models due to lack of logit access.

405
 406 **Model Performance Varies Across Alignment Dimensions.** No single model achieves optimal
 407 performance across all metrics. Llama-3.1 demonstrates the highest NVAS scores (75.75%), indi-
 408 cating strongest alignment with ground-truth MENA values, while frontier models GPT-4o-mini
 409 (89.47% CLCS) and Gemini (88.38% CLCS) exhibit superior cross-lingual consistency. Notably,
 410 regional specialist models (Fanar, ALLAM) do not outperform general-purpose models, suggesting
 411 that cultural alignment challenges persist despite targeted regional training.

412
 413 **Reasoning Prompts Can Reduce Cultural Alignment.** Across most settings, explicit reasoning
 414 prompts consistently decreased cultural alignment scores compared to zero-shot responses. This
 415 Reasoning-Induced Degradation phenomenon shows great decreases in NVAS scores for Mistral
 416 (-3.52%), Llama-3.1 (-6.96%), and Fanar (-6.12%). Our qualitative analysis identified three distinct
 417 failure modes underlying this phenomenon (see Appendix B.1).

418
 419 **Logit Leakage in Model Refusal Behavior.** Analysis of internal token probabilities reveals in-
 420 stances where models refuse to provide explicit answers while maintaining strong internal preferences
 421 for specific responses. Table 2 shows logit leakage rates ranging from 6.95% (ALLAM) to 47.50%
 422 (Fanar) in the with-reasoning condition. In these cases, models produce non-committal surface
 423 responses (e.g., "I cannot predict personal beliefs") while internal probability distributions show high
 424 confidence for particular answer choices (>75% probability mass). This discrepancy between surface
 425 responses and internal representations varies significantly across models, reasoning conditions, and
 426 language settings, as detailed in Appendix Table 3.

427 6.1 STRUCTURAL FAILURES IN CULTURAL REPRESENTATION

428 Principal Component Analysis of model responses reveals systematic patterns in how LLMs organize
 429 cultural representations across different conditions (Appendix E). These structural analyses demon-
 430 strate that observed inconsistencies reflect underlying representational failures rather than random
 431 variation, such as how reasoning conditions alter PCA clustering patterns across models, providing
 432 visual confirmation of the Reasoning effect.

432 Table 2: Logit Leakage Rate (%) by Model and Reasoning Condition.
433

434 Model	435 With-Reasoning (%)	436 Zero-Shot (%)
435 ALLaM	436 6.95	437 9.52
436 Fanar	437 47.50	438 33.65
437 Llama-3.1	438 20.97	439 5.86
438 Mistral	439 44.56	440 21.26
439 AYA	440 20.62	441 27.41

442 **Language-Based Clustering Overrides Cultural Distinctions.** Models demonstrate country-
443 specific differentiation when prompted in English, but this structure collapses when operating in
444 native languages. PCA analysis shows models cluster Arabic-speaking countries together while
445 isolating Persian-speaking Iran and Turkish-speaking Turkey, regardless of actual cultural similarities
446 between countries (Appendix E.3 and E.5).

447 **Models Maintain Distinct Cultural Positions.** When plotting model ‘neutral’ responses alongside
448 country-specific personas, the model’s own position consistently appears as an outlier, distant from
449 all MENA countries in the representational space (Appendix E.4), suggesting that the model’s own
450 beliefs differ substantially from the values of the MENA region.

452 7 DISCUSSION

454 **Implications for AI Alignment Theory.** The reasoning effect we observed challenges foundational
455 assumptions in AI safety research. If deliberative processes can change cultural alignment, this
456 suggests that current approaches have failed, as the model contains conflicts that, when engaged in
457 reasoning, alter its alignment. The scope of this effect remains unclear, we cannot determine whether
458 the degradation occurs across all cultural domains or only in value-laden questions concerning
459 underrepresented populations.

461 **Hidden Bias and the Limits of Surface-Level Safety.** The logit leakage phenomenon reveals
462 that current safety training may create a veneer of neutrality while preserving underlying biases
463 in model representations. This raises questions about what constitutes genuine alignment versus
464 performative compliance. If models maintain strong internal preferences while refusing to express
465 them, traditional evaluation methods that focus on outputs may systematically underestimate bias.
466 However, the interpretation of token probabilities as “beliefs” remains contested, and we cannot
467 definitively claim that these internal states represent conscious biases.

468 **Multilingual Training and Cultural Essentialism.** The collapse into language-based clustering
469 suggests that current multilingual training approaches may inadvertently promote linguistic essen-
470 tialism, treating language as a perfect proxy for culture. This has serious implications for global AI
471 deployment, as it implies models may homogenize diverse cultural contexts within language families
472 while artificially amplifying differences between them.

474 8 CONCLUSION

476 In this paper, we introduce MENAValues, a comprehensive, empirically-grounded benchmark for
477 assessing the cultural alignment of LLMs with populations from the Middle East and North Africa.
478 Our evaluation of seven diverse models reveals significant misalignments and inconsistencies, which
479 are highly sensitive to prompt language, perspective, and reasoning. We identify and analyze several
480 critical phenomena which highlight fundamental challenges for global AI alignment and safety.
481 As LLMs are deployed globally, ensuring they can respectfully and accurately navigate diverse
482 cultural contexts is paramount. The MENAValues benchmark and our analytical framework provide a
483 robust methodology for diagnosing these complex alignment failures. Future work should focus on
484 developing methods to improve cross-cultural and cross-lingual consistency and expand such deep
485 evaluations to other underrepresented regions, paving the way for AI that is truly aligned with the
rich diversity of human values.

486 REFERENCES
487488 Abubakar Abid, Maheen Farooqi, and James Zou. Persistent anti-muslim bias in large language
489 models, 2021. URL <https://arxiv.org/abs/2101.05783>.490 Muhammad Farid Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Shivdutt Singh,
491 Alham Fikri Aji, Jacki O'Neill, Ashutosh Modi, and Monojit Choudhury. Towards measuring
492 and modeling “culture” in LLMs: A survey. In Yaser Al-Onaizan, Mohit Bansal, and
493 Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural
494 Language Processing*, pp. 15763–15784, Miami, Florida, USA, November 2024. Association
495 for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.882. URL [https://aclanthology.org/2024.emnlp-main.882/](https://aclanthology.org/2024.emnlp-main.882).496
497 Badr AlKhamissi, Muhammad ElNokrashy, Mai AlKhamissi, and Mona Diab. Investigating
498 cultural alignment of large language models. In Lun-Wei Ku, Andre Martins, and Vivek
499 Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational
500 Linguistics (Volume 1: Long Papers)*, pp. 12404–12422, Bangkok, Thailand, August
501 2024a. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.671. URL
502 [https://aclanthology.org/2024.acl-long.671/](https://aclanthology.org/2024.acl-long.671).503
504 Badr AlKhamissi, Muhammad ElNokrashy, Mai AlKhamissi, and Mona Diab. Investigating cul-
505 tural alignment of large language models, 2024b. URL <https://arxiv.org/abs/2402.13231>.506
507 Arab Center for Research and Policy Studies. Arab opinion index 2022. <https://arabindex.dohainstitute.org/EN/Pages/Arab-Opinion-Index-2022.aspx>, 2022. Eighth
508 wave of the Arab Opinion Index, based on face-to-face interviews with 33,300 respondents across
509 14 Arab countries.510
511 Arnav Arora, Lucie-aimée Kaffee, and Isabelle Augenstein. Probing pre-trained language models
512 for cross-cultural differences in values. In Sunipa Dev, Vinodkumar Prabhakaran, David Ife-
513 oluwa Adelani, Dirk Hovy, and Luciana Benotti (eds.), *Proceedings of the First Workshop
514 on Cross-Cultural Considerations in NLP (C3NLP)*, pp. 114–130, Dubrovnik, Croatia, May
515 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.c3nlp-1.12. URL
516 <https://aclanthology.org/2023.c3nlp-1.12/>.517
518 Nishant Balepur, Abhilasha Ravichander, and Rachel Rudinger. Artifacts or abduction: How do
519 LLMs answer multiple-choice questions without the question? In Lun-Wei Ku, Andre Martins,
520 and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for
521 Computational Linguistics (Volume 1: Long Papers)*, pp. 10308–10330, Bangkok, Thailand,
522 August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.555.
523 URL <https://aclanthology.org/2024.acl-long.555/>.524
525 Yong Cao, Li Zhou, Seolhwa Lee, Laura Cabello, Min Chen, and Daniel Herscovitch. Assessing
526 cross-cultural alignment between ChatGPT and human societies: An empirical study. In Sunipa
527 Dev, Vinodkumar Prabhakaran, David Ifeoluwa Adelani, Dirk Hovy, and Luciana Benotti (eds.),
528 *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pp. 53–67,
529 Dubrovnik, Croatia, May 2023. Association for Computational Linguistics. doi: 10.18653/v1/
2023.c3nlp-1.7. URL <https://aclanthology.org/2023.c3nlp-1.7/>.530
531 Myra Cheng, Esin Durmus, and Dan Jurafsky. Marked personas: Using natural language
532 prompts to measure stereotypes in language models. In Anna Rogers, Jordan Boyd-Graber,
533 and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for
534 Computational Linguistics (Volume 1: Long Papers)*, pp. 1504–1532, Toronto, Canada, July
535 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.84. URL
<https://aclanthology.org/2023.acl-long.84/>.536
537 Myra Cheng, Kristina Gligoric, Tiziano Piccardi, and Dan Jurafsky. AnthroScore: A computational
538 linguistic measure of anthropomorphism. In Yvette Graham and Matthew Purver (eds.), *Proceed-
539 ings of the 18th Conference of the European Chapter of the Association for Computational Linguis-
tics (Volume 1: Long Papers)*, pp. 807–825, St. Julian’s, Malta, March 2024. Association for Com-
putational Linguistics. URL <https://aclanthology.org/2024.eacl-long.49/>.

540 John Dang, Shivalika Singh, Daniel D’souza, Arash Ahmadian, Alejandro Salamanca, Madeline
 541 Smith, Aidan Peppin, Sungjin Hong, Manoj Govindassamy, Terrence Zhao, Sandra Kublik, Meor
 542 Amer, Viraat Aryabumi, Jon Ander Campos, Yi-Chern Tan, Tom Kocmi, Florian Strub, Nathan
 543 Grinsztajn, Yannis Flet-Berliac, Acyr Locatelli, Hangyu Lin, Dwarak Talupuru, Bharat Venkitesh,
 544 David Cairuz, Bowen Yang, Tim Chung, Wei-Yin Ko, Sylvie Shang Shi, Amir Shukayev, Sammie
 545 Bae, Aleksandra Piktus, Roman Castagné, Felipe Cruz-Salinas, Eddie Kim, Lucas Crawhall-Stein,
 546 Adrien Morisot, Sudip Roy, Phil Blunsom, Ivan Zhang, Aidan Gomez, Nick Frosst, Marzieh Fadaee,
 547 Beyza Ermis, Ahmet Üstün, and Sara Hooker. Aya expanse: Combining research breakthroughs
 548 for a new multilingual frontier, 2024. URL <https://arxiv.org/abs/2412.04261>.

549 Ameet Deshpande, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, and Karthik Narasimhan.
 550 Toxicity in chatgpt: Analyzing persona-assigned language models. In Houda Bouamor, Juan
 551 Pino, and Kalika Bali (eds.), *Findings of the Association for Computational Linguistics: EMNLP*
 552 2023, pp. 1236–1270, Singapore, December 2023. Association for Computational Linguistics.
 553 doi: 10.18653/v1/2023.findings-emnlp.88. URL <https://aclanthology.org/2023.findings-emnlp.88/>.

555 Awantee Deshpande, Dana Ruiter, Marius Mosbach, and Dietrich Klakow. StereoKG: Data-driven
 556 knowledge graph construction for cultural knowledge and stereotypes. In Kanika Narang, Aida
 557 Mostafazadeh Davani, Lambert Mathias, Bertie Vidgen, and Zeerak Talat (eds.), *Proceedings of the*
 558 *Sixth Workshop on Online Abuse and Harms (WOAH)*, pp. 67–78, Seattle, Washington (Hybrid),
 559 July 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.woah-1.7. URL
 560 <https://aclanthology.org/2022.woah-1.7/>.

561 Esin Durmus, Karina Nguyen, Thomas I. Liao, Nicholas Schiefer, Amanda Askell, Anton Bakhtin,
 562 Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, Liane Lovitt, Sam McCandlish,
 563 Orowa Sikder, Alex Tamkin, Janel Thamkul, Jared Kaplan, Jack Clark, and Deep Ganguli.
 564 Towards measuring the representation of subjective global opinions in language models, 2024a.
 565 URL <https://arxiv.org/abs/2306.16388>.

566 Esin Durmus, Karina Nguyen, Thomas I. Liao, Nicholas Schiefer, Amanda Askell, Anton Bakhtin,
 567 Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, Liane Lovitt, Sam McCandlish,
 568 Orowa Sikder, Alex Tamkin, Janel Thamkul, Jared Kaplan, Jack Clark, and Deep Ganguli.
 569 Towards measuring the representation of subjective global opinions in language models, 2024b.
 570 URL <https://arxiv.org/abs/2306.16388>.

572 Julen Etxaniz, Gorka Azkune, Aitor Soroa, Oier Lopez de Lacalle, and Mikel Artetxe. Do multilingual
 573 language models think better in English? In Kevin Duh, Helena Gomez, and Steven Bethard
 574 (eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for*
 575 *Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pp. 550–564,
 576 Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/
 577 2024.naacl-short.46. URL <https://aclanthology.org/2024.naacl-short.46/>.

578 Negar Foroutan, Mohammadreza Banaei, Rémi Lebret, Antoine Bosselut, and Karl Aberer. Discovering
 579 language-neutral sub-networks in multilingual language models. In Yoav Goldberg, Zornitsa
 580 Kozareva, and Yue Zhang (eds.), *Proceedings of the 2022 Conference on Empirical Methods in*
 581 *Natural Language Processing*, pp. 7560–7575, Abu Dhabi, United Arab Emirates, December
 582 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.513. URL
 583 [https://aclanthology.org/2022.emnlp-main.513/](https://aclanthology.org/2022.emnlp-main.513).

584 Yi Fung, Ruining Zhao, Jae Doo, Chenkai Sun, and Heng Ji. Massively multi-cultural knowledge
 585 acquisition & lm benchmarking, 2024. URL <https://arxiv.org/abs/2402.09369>.

586 Iason Gabriel and Vafa Ghazavi. The challenge of value alignment: from fairer algorithms to ai
 587 safety, 2021. URL <https://arxiv.org/abs/2101.06060>.

589 Google Cloud. Gemini 2.5 flash-lite. <https://cloud.google.com/vertex-ai/generative-ai/docs/models/gemini/2-5-flash-lite>, 2025. Accessed: 2025-08-24.

593 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
 594 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan,

594 Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev,
 595 Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru,
 596 Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak,
 597 Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu,
 598 Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle
 599 Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego
 600 Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova,
 601 Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel
 602 Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon,
 603 Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan
 604 Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet,
 605 Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde,
 606 Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie
 607 Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua
 608 Saxe, Junteng Jia, Kalyan Vasudevan Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak,
 609 Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley
 610 Chiu, Kunal Bhalla, Kushal Lakhota, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence
 611 Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas
 612 Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri,
 613 Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie
 614 Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes
 615 Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne,
 616 Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajwal
 617 Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong,
 618 Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic,
 619 Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie
 620 Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana
 621 Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie,
 622 Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon
 623 Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan,
 624 Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas
 625 Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami,
 626 Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vojeti,
 627 Vítor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier
 628 Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao
 629 Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song,
 630 Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe
 631 Papakipos, Aditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya
 632 Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenber, Alexei
 633 Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu,
 634 Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit
 635 Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury,
 636 Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer,
 637 Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu,
 638 Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido,
 639 Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu
 640 Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer,
 641 Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu,
 642 Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc
 643 Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily
 644 Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers,
 645 Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank
 646 Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee,
 647 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,

648 Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai
 649 Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik
 650 Veeraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle
 651 Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng
 652 Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish
 653 Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim
 654 Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle
 655 Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang,
 656 Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam,
 657 Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier,
 658 Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia
 659 Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro
 660 Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani,
 661 Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy,
 662 Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin
 663 Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu,
 664 Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh
 665 Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay,
 666 Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang,
 667 Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie
 668 Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta,
 669 Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman,
 670 Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun
 671 Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria
 672 Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru,
 673 Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz,
 674 Will Constable, Xiaocheng Tang, Xiaoqian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv
 675 Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi,
 676 Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait,
 677 Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. The
 678 llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.

679 Vipul Gupta, Pranav Narayanan Venkit, Shomir Wilson, and Rebecca Passonneau. Sociodemographic
 680 bias in language models: A survey and forward path. In Agnieszka Faleńska, Christine Basta,
 681 Marta Costa-jussà, Seraphina Goldfarb-Tarrant, and Debora Nozza (eds.), *Proceedings of the 5th
 682 Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pp. 295–322, Bangkok,
 Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.
 683 gebnlp-1.19. URL <https://aclanthology.org/2024.gebnlp-1.19/>.

684 Christian Haerpfer, Ronald Inglehart, Alejandro Moreno, Christian Welzel, Kseniya Kizilova, Juan
 685 Diez-Medrano, Marta Lagos, Pippa Norris, Eduard Ponarin, Bi Puranen, et al. World values survey:
 686 Round seven – country-pooled datafile version 6.0. <https://doi.org/10.14281/18241.24>, 2022. Editors.

687 Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob
 688 Steinhardt. Aligning ai with shared human values, 2023. URL <https://arxiv.org/abs/2008.02275>.

689 Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie
 690 Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza
 691 Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. Challenges and strategies in cross-
 692 cultural NLP. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Proceedings of
 693 the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*,
 694 pp. 6997–7013, Dublin, Ireland, May 2022a. Association for Computational Linguistics. doi:
 695 10.18653/v1/2022.acl-long.482. URL [https://aclanthology.org/2022.acl-long.482/](https://aclanthology.org/2022.acl-long.482).

696 Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie
 697 Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza
 698 Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. Challenges and strategies in cross-
 699 cultural NLP. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Proceedings of
 700 the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*,
 701 pp. 6997–7013, Dublin, Ireland, May 2022a. Association for Computational Linguistics. doi:
 10.18653/v1/2022.acl-long.482. URL [https://aclanthology.org/2022.acl-long.482/](https://aclanthology.org/2022.acl-long.482).

702 *the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*,
 703 pp. 6997–7013, Dublin, Ireland, May 2022b. Association for Computational Linguistics. doi:
 704 10.18653/v1/2022.acl-long.482. URL <https://aclanthology.org/2022.acl-long.482/>.

705

706 Huang Huang, Fei Yu, Jianqing Zhu, Xuening Sun, Hao Cheng, Song Dingjie, Zhihong Chen, Mosen
 707 Alharthi, Bang An, Juncai He, Ziche Liu, Junying Chen, Jianquan Li, Benyou Wang, Lian Zhang,
 708 Ruoyu Sun, Xiang Wan, Haizhou Li, and Jinchao Xu. AceGPT, localizing large language models
 709 in Arabic. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024*
 710 *Conference of the North American Chapter of the Association for Computational Linguistics: Human*
 711 *Language Technologies (Volume 1: Long Papers)*, pp. 8139–8163, Mexico City, Mexico,
 712 June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.450.
 713 URL <https://aclanthology.org/2024.naacl-long.450/>.

714

715 Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke Zettlemoyer, Hannaneh
 716 Hajishirzi, Yejin Choi, and Prithviraj Ammanabrolu. Personalized soups: Personalized large
 717 language model alignment via post-hoc parameter merging, 2023. URL <https://arxiv.org/abs/2310.11564>.

718

719 Akshita Jha, Aida Mostafazadeh Davani, Chandan K Reddy, Shachi Dave, Vinodkumar Prabhakaran,
 720 and Sunipa Dev. SeeGULL: A stereotype benchmark with broad geo-cultural coverage leveraging
 721 generative models. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings*
 722 *of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*,
 723 pp. 9851–9870, Toronto, Canada, July 2023. Association for Computational Linguistics. doi:
 724 10.18653/v1/2023.acl-long.548. URL [https://aclanthology.org/2023.acl-long.548/](https://aclanthology.org/2023.acl-long.548).

725

726 Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot,
 727 Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier,
 728 Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas
 729 Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023. URL <https://arxiv.org/abs/2310.06825>.

730

731 Nitish Joshi, Javier Rando, Abulhair Saparov, Najoung Kim, and He He. Personas as a way to
 732 model truthfulness in language models. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung
 733 Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language*
 734 *Processing*, pp. 6346–6359, Miami, Florida, USA, November 2024. Association for Computational
 735 Linguistics. doi: 10.18653/v1/2024.emnlp-main.364. URL [https://aclanthology.org/2024.emnlp-main.364/](https://aclanthology.org/2024.emnlp-main.364).

736

737 Julia Kharchenko, Tanya Roosta, Aman Chadha, and Chirag Shah. How well do llms represent values
 738 across cultures? empirical analysis of llm responses based on hofstede cultural dimensions, 2024.
 739 URL <https://arxiv.org/abs/2406.14805>.

740

741 Hannah Rose Kirk, Bertie Vidgen, Paul Röttger, and Scott A. Hale. Personalisation within bounds: A
 742 risk taxonomy and policy framework for the alignment of large language models with personalised
 743 feedback, 2023. URL <https://arxiv.org/abs/2303.05453>.

744

745 Hannah Rose Kirk, Alexander Whitefield, Paul Röttger, Andrew Bean, Katerina Margatina, Juan Ciro,
 746 Rafael Mosquera, Max Bartolo, Adina Williams, He He, Bertie Vidgen, and Scott A. Hale. The
 747 prism alignment dataset: What participatory, representative and individualised human feedback
 748 reveals about the subjective and multicultural alignment of large language models, 2024. URL
 749 <https://arxiv.org/abs/2404.16019>.

750

751 Cheng Li, Mengzhou Chen, Jindong Wang, Sunayana Sitaram, and Xing Xie. Culturellm: Incorporat-
 752 ing cultural differences into large language models, 2024a. URL <https://arxiv.org/abs/2402.10946>.

753

754 Jialin Li, Junli Wang, Junjie Hu, and Ming Jiang. How well do llms identify cultural unity in
 755 diversity?, 2024b. URL <https://arxiv.org/abs/2408.05102>.

756 Wangyue Li, Liangzhi Li, Tong Xiang, Xiao Liu, Wei Deng, and Noa Garcia. Can multiple-choice
 757 questions really be useful in detecting the abilities of LLMs? In Nicoletta Calzolari, Min-Yen
 758 Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (eds.), *Proceedings of
 759 the 2024 Joint International Conference on Computational Linguistics, Language Resources and
 760 Evaluation (LREC-COLING 2024)*, pp. 2819–2834, Torino, Italia, May 2024c. ELRA and ICCL.
 761 URL <https://aclanthology.org/2024.lrec-main.251/>.

762 Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. Towards understand-
 763 ing and mitigating social biases in language models, 2021. URL <https://arxiv.org/abs/2106.13219>.

764

765 Margaret Mitchell, Giuseppe Attanasio, Ioana Baldini, Miruna Clinciu, Jordan Clive, Pieter Delobelle,
 766 Manan Dey, Sil Hamilton, Timm Dill, Jad Doughman, Ritam Dutt, Avijit Ghosh, Jessica Zosa
 767 Forde, Carolin Holtermann, Lucie-Aimée Kaffee, Tanmay Laud, Anne Lauscher, Roberto L Lopez-
 768 Davila, Maraim Masoud, Nikita Nangia, Anaelia Ovalle, Giada Pistilli, Dragomir Radev, Beatrice
 769 Savoldi, Vipul Raheja, Jeremy Qin, Esther Ploeger, Arjun Subramonian, Kaustubh Dhole, Kaiser
 770 Sun, Amirbek Djanibekov, Jonibek Mansurov, Kayo Yin, Emilio Villa Cueva, Sagnik Mukherjee,
 771 Jerry Huang, Xudong Shen, Jay Gala, Hamdan Al-Ali, Tair Djanibekov, Nurdaulet Mukhituly,
 772 Shangrui Nie, Shanya Sharma, Karolina Stanczak, Eliza Szczechla, Tiago Timponi Torrent,
 773 Deepak Tunuguntla, Marcelo Viridiano, Oskar Van Der Wal, Adina Yakefu, Aurélie Névéol,
 774 Mike Zhang, Sydney Zink, and Zeerak Talat. SHADES: Towards a multilingual assessment
 775 of stereotypes in large language models. In Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.),
 776 *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for
 777 Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 11995–
 778 12041, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. ISBN
 779 979-8-89176-189-6. URL <https://aclanthology.org/2025.naacl-long.600/>.

780 Moin Nadeem, Anna Bethke, and Siva Reddy. StereoSet: Measuring stereotypical bias in pretrained
 781 language models. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Proceed-
 782 ings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th
 783 International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp.
 784 5356–5371, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/
 785 2021.acl-long.416. URL <https://aclanthology.org/2021.acl-long.416/>.

786 Tarek Naous, Michael J. Ryan, Alan Ritter, and Wei Xu. Having beer after prayer? measuring cultural
 787 bias in large language models, 2024a. URL <https://arxiv.org/abs/2305.14456>.

788

789 Tarek Naous, Michael J. Ryan, Alan Ritter, and Wei Xu. Having beer after prayer? measuring cultural
 790 bias in large language models, 2024b. URL <https://arxiv.org/abs/2305.14456>.

791 Tarek Naous, Michael J. Ryan, Alan Ritter, and Wei Xu. Having beer after prayer? mea-
 792 suring cultural bias in large language models. In Lun-Wei Ku, Andre Martins, and Vivek
 793 Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Compu-
 794 tational Linguistics (Volume 1: Long Papers)*, pp. 16366–16393, Bangkok, Thailand, August
 795 2024c. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.862. URL
 796 <https://aclanthology.org/2024.acl-long.862/>.

797 OpenAI. Gpt-4o mini: advancing cost-efficient intelligence. <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>, 2025. Accessed:
 798 2025-08-24.

799

800 Jonas Pfeiffer, Naman Goyal, Xi Lin, Xian Li, James Cross, Sebastian Riedel, and Mikel Artetxe.
 801 Lifting the curse of multilinguality by pre-training modular transformers. In Marine Carpuat, Marie-
 802 Catherine de Marneffe, and Ivan Vladimir Meza Ruiz (eds.), *Proceedings of the 2022 Conference of
 803 the North American Chapter of the Association for Computational Linguistics: Human Language
 804 Technologies*, pp. 3479–3495, Seattle, United States, July 2022. Association for Computational
 805 Linguistics. doi: 10.18653/v1/2022.naacl-main.255. URL <https://aclanthology.org/2022.naacl-main.255/>.

806

807 Jirui Qi, Raquel Fernández, and Arianna Bisazza. Cross-lingual consistency of factual knowledge in
 808 multilingual language models. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings
 809 of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 10650–10666,

810 Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.
 811 emnlp-main.658. URL <https://aclanthology.org/2023.emnlp-main.658/>.
 812

813 Mohammed Al-Maghribi Research. Allam-thinking: Arabic large language model with enhanced rea-
 814 soning capabilities. <https://huggingface.co/almaghrabima/ALLaM-Thinking>,
 815 2025.

816 Michael J Ryan, William Held, and Diyi Yang. Unintended impacts of LLM alignment on global
 817 representation. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd*
 818 *Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp.
 819 16121–16140, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi:
 820 10.18653/v1/2024.acl-long.853. URL <https://aclanthology.org/2024.acl-long.853/>.
 821

822 Lingfeng Shen, Weiting Tan, Sihao Chen, Yunmo Chen, Jingyu Zhang, Haoran Xu, Boyuan Zheng,
 823 Philipp Koehn, and Daniel Khashabi. The language barrier: Dissecting safety challenges of LLMs
 824 in multilingual contexts. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of*
 825 *the Association for Computational Linguistics: ACL 2024*, pp. 2668–2680, Bangkok, Thailand,
 826 August 2024a. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.156.
 827 URL <https://aclanthology.org/2024.findings-acl.156/>.
 828

829 Siqi Shen, Lajanugen Logeswaran, Moontae Lee, Honglak Lee, Soujanya Poria, and Rada Mi-
 830 halcea. Understanding the capabilities and limitations of large language models for cultural
 831 commonsense. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the*
 832 *2024 Conference of the North American Chapter of the Association for Computational Linguistics:*
 833 *Human Language Technologies (Volume 1: Long Papers)*, pp. 5668–5680, Mexico City, Mexico,
 834 June 2024b. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.316.
 835 URL <https://aclanthology.org/2024.naacl-long.316/>.
 836

836 Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. Societal biases in language
 837 generation: Progress and challenges. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli
 838 (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics*
 839 *and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long*
 840 *Papers)*, pp. 4275–4293, Online, August 2021. Association for Computational Linguistics. doi:
 841 10.18653/v1/2021.acl-long.330. URL <https://aclanthology.org/2021.acl-long.330/>.
 842

843 Weiyan Shi, Ryan Li, Yutong Zhang, Caleb Ziems, Chunhua Yu, Raya Horesh, Rogério Abreu
 844 de Paula, and Diyi Yang. Culturebank: An online community-driven knowledge base towards cul-
 845 turally aware language technologies, 2024. URL <https://arxiv.org/abs/2404.15238>.
 846

847 Yan Tao, Olga Viberg, Ryan S Baker, and René F Kizilcec. Cultural bias and cultural alignment of
 848 large language models. *PNAS Nexus*, 3(9), September 2024a. ISSN 2752-6542. doi: 10.1093/
 849 pnasnexus/pgae346. URL <http://dx.doi.org/10.1093/pnasnexus/pgae346>.
 850

850 Yan Tao, Olga Viberg, Ryan S Baker, and René F Kizilcec. Cultural bias and cultural alignment of
 851 large language models. *PNAS Nexus*, 3(9), September 2024b. ISSN 2752-6542. doi: 10.1093/
 852 pnasnexus/pgae346. URL <http://dx.doi.org/10.1093/pnasnexus/pgae346>.
 853

854 Fanar Team, Ummar Abbas, Mohammad Shahmeer Ahmad, Firoj Alam, Enes Altinisik, Ehsannedin
 855 Asgari, Yazan Boshmaf, Sabri Boughorbel, Sanjay Chawla, Shammur Chowdhury, Fahim Dalvi,
 856 Kareem Darwish, Nadir Durrani, Mohamed Elfeky, Ahmed Elmagarmid, Mohamed Eltabakh,
 857 Masoomali Fatehkia, Anastasios Fragkopoulos, Maram Hasanain, Majd Hawasly, Mus'ab Husaini,
 858 Soon-Gyo Jung, Ji Kim Lucas, Walid Magdy, Safa Messaoud, Abubakr Mohamed, Tasnim
 859 Mohiuddin, Basel Mousi, Hamdy Mubarak, Ahmad Musleh, Zan Naeem, Mourad Ouzzani,
 860 Dorde Popovic, Amin Sadeghi, Husrev Taha Sencar, Mohammed Shinoy, Omar Sinan, Yifan
 861 Zhang, Ahmed Ali, Yassine El Kheir, Xiaosong Ma, and Chaoyi Ruan. Fanar: An arabic-centric
 862 multimodal generative ai platform, 2025. URL <https://arxiv.org/abs/2501.13944>.
 863

Bin Wang, Zhengyuan Liu, Xin Huang, Fangkai Jiao, Yang Ding, AiTi Aw, and Nancy Chen.
 SeaEval for multilingual foundation models: From cross-lingual alignment to cultural reasoning.

864 In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Conference*
865 *of the North American Chapter of the Association for Computational Linguistics: Human*
866 *Language Technologies (Volume 1: Long Papers)*, pp. 370–390, Mexico City, Mexico, June
867 2024a. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.22. URL
868 <https://aclanthology.org/2024.naacl-long.22/>.

869 Wenzuan Wang, Wenxiang Jiao, Jingyuan Huang, Ruyi Dai, Jen-tse Huang, Zhaopeng Tu, and
870 Michael Lyu. Not all countries celebrate thanksgiving: On the cultural dominance in large
871 language models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the*
872 *62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*,
873 pp. 6349–6384, Bangkok, Thailand, August 2024b. Association for Computational Linguistics. doi:
874 10.18653/v1/2024.acl-long.345. URL [https://aclanthology.org/2024.acl-long.](https://aclanthology.org/2024.acl-long.345)
875 345/.

876 Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. Large language models are
877 not robust multiple choice selectors, 2024. URL <https://arxiv.org/abs/2309.03882>.

878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

918 LIMITATIONS
919

920 While the MENAVales Benchmark provides valuable insights into cultural alignment, several
921 limitations merit consideration. First, our benchmark relies on survey data that, despite rigorous
922 methodologies, may not capture the full complexity of values within MENA societies. We also
923 acknowledge the general limitations of using multiple-choice questionnaires to measure complex
924 traits, though our methodology (using probability distributions, multiple framings, and consistency
925 checks) is designed to mitigate these concerns. Second, our analysis focuses on seven models, and
926 findings may not generalize to all architectures or scales. Third, translations between English and
927 native languages, though validated by humans, may introduce subtle semantic shifts.

928 Our token-level probability analysis approach has several limitations worth noting. First, the method
929 relies on identifying the correct token position for extracting answer probabilities, which can be
930 challenging when models produce unexpected response patterns. Second, our approach focuses on the
931 first few tokens of generation, which may not capture the full deliberative process in more complex
932 responses. The normalization procedure we employ, while necessary for comparing probabilities
933 across different answer options, can sometimes amplify small differences in low-probability scenarios.
934 Additionally, our implementation analyzes only the top token candidates for each answer option,
935 potentially overlooking complex tokenization patterns where answers might be split across multiple
936 tokens or represented through unexpected encodings. We encourage more sophisticated approaches
937 to involve analyzing the probability flow across entire generation sequences.

938 Looking forward, the phenomena identified in this study, particularly Reasoning-Induced Degradation,
939 Logit Leakage, and Cross-Lingual Value Shifts, represent critical avenues for future research. A
940 deeper investigation into their underlying causal mechanisms is necessary, not only to better explain
941 these complex behaviors but also to develop effective mitigation strategies. Ultimately, understanding
942 and addressing these issues will be essential for building LLMs that are more transparent, reliable,
943 and genuinely aligned with the diverse spectrum of human cultures.

944
945 ETHICS STATEMENT
946

947 Our research introduces the MENAVales benchmark as a diagnostic tool to identify and measure
948 cultural misalignment and biases in LLMs with respect to the MENA region. Our work is intended
949 to foster the development of more culturally aware and inclusive AI, not to build models that make
950 prescriptive judgments about MENA societies or to reinforce stereotypes.

951 The ground-truth data in our benchmark is derived from the publicly available, anonymized World
952 Values Survey and Arab Opinion Index. It is essential to recognize that this data is *descriptive*,
953 reflecting the reported opinions of survey respondents, and not *prescriptive* of how individuals or
954 societies should behave. While these surveys employ rigorous methodologies, we acknowledge that
955 no dataset can fully capture the complexity and diversity of the MENA region. We have made efforts
956 to preserve this diversity by including 16 countries and analyzing full human response distributions.

957 A high alignment score (NVAS) in our benchmark indicates that a model’s output is closer to the
958 majority human opinion from the survey data. However, a high score is not inherently “better” or
959 more desirable. It could mean the model is accurately reflecting benign cultural norms, but it could
960 also mean the model is successfully reproducing harmful societal biases (e.g., regarding gender roles
961 or stereotypes) that may be present in the human data. We emphasize that our results should be
962 interpreted as diagnostic signals of a model’s underlying value system, not as a prescriptive target for
963 alignment.

964 The complexity of normative alignment raises questions about when and whether AI models should
965 align with population majorities versus other normative frameworks. We do not advocate for blind
966 adherence to majority opinion, which can perpetuate discrimination against minority groups or
967 vulnerable populations. Rather, our benchmark serves a diagnostic function: revealing how models
968 currently represent cultural values and where systematic biases occur. The question of which values
969 AI systems should embody is a complex societal decision that goes beyond technical evaluation.
970 Different stakeholders may reasonably prioritize different alignment targets, such as local cultural
971 authenticity, universal human rights principles, or context-dependent balancing of these concerns.
Our benchmark makes visible the tradeoffs involved in these choices rather than resolving them. For

972 example, high alignment with local gender role attitudes might reflect cultural sensitivity in some
 973 contexts while contradicting universal equality principles in others. We believe transparency about
 974 these tensions is preferable to implicit bias toward any single normative framework.

975 The phenomena we identify, particularly Logit Leakage, raise significant concerns for AI safety and
 976 transparency. This suggests that current alignment techniques may be insufficient, merely teaching
 977 models to hide their biases rather than resolving them. Our work underscores the need for deeper,
 978 more fundamental approaches to AI safety that go beyond surface-level outputs and engage seriously
 979 with the normative complexity of cross-cultural deployment.

980 Survey data and multiple-choice formats cannot capture the full complexity of human values. More-
 981 over, the MENA region is not a monolith, and our benchmark should not be used to essentialize or
 982 over-generalize the diverse beliefs of this region. We encourage practitioners to view our benchmark
 983 as a starting point for identifying cultural misalignment, with the hope that future work will expand
 984 this type of deep evaluation to other underrepresented regions while continuing to grapple with the
 985 fundamental normative questions that cross-cultural AI deployment raises.

988 REPRODUCIBILITY STATEMENT

990 We have made substantial efforts to ensure the reproducibility of this work. Our MENAValues
 991 benchmark dataset, constructed from publicly available World Values Survey Wave 7 and Arab
 992 Opinion Index 2022 data, will be made available upon publication along with our complete LLM
 993 evaluation outputs.

994 Our evaluation framework is thoroughly documented in Section 4, including mathematical formu-
 995 lations for all metrics (NVAS, FCS, CLCS, SPD) and our token probability analysis methodology.
 996 The complete experimental setup, including model configurations and prompting templates across
 997 all three perspective framings (neutral, persona, observer) and languages (English, Arabic, Persian,
 998 Turkish) is detailed in Section 4 and will be available on our GitHub. Our code for conducting the
 999 evaluation, including logit extraction procedures, statistical analysis, and PCA visualizations, will be
 1000 released as supplementary materials.

1001 The substantial scale of our evaluation (over 820,000 data points across 7 models, 16 countries, 864
 1002 questions, multiple conditions) and our approach to documenting experimental procedures should
 1003 enable full replication of our results.

1006 USE OF LARGE LANGUAGE MODELS

1008 Large language models were used in limited capacity during this research as general-purpose as-
 1009 sistance tools during this research in two capacities: (1) writing assistance for improving clarity,
 1010 grammar, and organization of the manuscript text, and (2) code generation and debugging assis-
 1011 tance for data processing and visualization scripts. LLMs were not involved in research design,
 1012 methodology development, interpretation of results, or generation of core research ideas and contribu-
 1013 tions. All substantive content, including the research framework, experimental design, and scientific
 1014 conclusions, was developed entirely by the human authors.

1017 A BENCHMARK CURATION DETAILS

1020 A.1 QUESTION SELECTION CRITERIA

1022 Our process for selecting the 864 questions from the source surveys (WVS-7 and AOI-2022) involved
 1023 manual validation to ensure each question was value-centric. We filtered out questions that were
 1024 purely factual (e.g., "Which of these organizations do you belong to?"), demographic (e.g., "What is
 1025 your age?"), or otherwise irrelevant to capturing beliefs or attitudes. This focused curation ensures
 the benchmark is concentrated on assessing cultural and social values.

1026 **B QUALITATIVE ANALYSIS OF MODEL BEHAVIORS**
10271028 **B.1 REASONING-INDUCED PERFORMANCE DEGRADATION: QUALITATIVE ANALYSIS**
10291030 As noted in the main paper, prompting LLMs to provide reasoning often degrades their cultural
1031 alignment. Our qualitative analysis of model outputs identified three distinct failure modes that
1032 explain this phenomenon:
10331034

- 1035 **1. Cultural Stereotyping and Overgeneralization:** When asked to reason, models often fall
1036 back on broad, often Western-centric, stereotypes about the MENA region. They fail to
1037 capture the nuanced diversity within and across different societies, producing rationales that
1038 treat "MENA" or a specific nationality as a monolith.
- 1039 **2. Cultural Value Projection:** The reasoning process appears to activate the model's under-
1040 lying, predominantly Western-liberal value system. Models often generate justifications
1041 that align with Western norms (e.g., prioritizing individual autonomy or secularism) even
1042 if those justifications lead to a final answer that conflicts with the empirically documented
1043 local values.
- 1044 **3. Safety-Induced Self-Censorship:** The request for reasoning on a potentially sensitive
1045 cultural topic frequently triggers overly cautious behavior. This leads to hedged, vague, or
1046 generic responses that avoid taking a culturally specific stance. For instance, a model might
1047 deflect by stating, "As an AI, I cannot have personal beliefs," or provide a generic response
1048 like, "This is a complex issue with diverse viewpoints," effectively failing to answer the
1049 question from the requested cultural perspective.

1050 These failure modes suggest that the intuitive approach of "making models think harder" can be
1051 counterproductive for culturally-nuanced tasks, as the reasoning process itself can introduce or
1052 amplify biases.
10531054 **C REGIONAL HETEROGENEITY ANALYSIS: UNDERSTANDING VARIANCE
1055 WITHIN MENA COUNTRIES**
10561058 The Middle East and North Africa region is frequently conceptualized as a monolithic entity in
1059 cross-cultural research, yet it comprises a deeply heterogeneous collection of nations with distinct
1060 historical trajectories, political systems, and socio-cultural fabrics. To better contextualize our
1061 benchmark findings and validate the representativeness of our cultural alignment evaluation, we
1062 present a comprehensive quantitative analysis of regional similarities and differences using the same
1063 foundational datasets that underpin MENAValues.1064 This analysis serves two critical purposes: (1) it demonstrates the empirical basis for treating MENA
1065 as a coherent yet internally diverse cultural region, and (2) it provides insights into which dimensions
1066 of cultural values show convergence versus divergence across the region, informing the interpretation
1067 of our LLM alignment results.
10681069 **C.1 WORLD VALUES SURVEY ANALYSIS**
10701071 We conducted PCA on the WVS-7 subset comprising 9 countries with available data: Egypt, Iran,
1072 Iraq, Jordan, Lebanon, Libya, Morocco, Turkey, and Tunisia. This sample notably includes major
1073 non-Arab states (Iran and Turkey), providing a unique lens for examining regional dynamics beyond
1074 the Arab-non-Arab dichotomy.
10751076 **C.1.1 PRINCIPAL COMPONENT STRUCTURE**
10771078 The first two principal components collectively account for 42.53% of the total variance, with PC1
1079 explaining 27.21% and PC2 explaining 15.32%. The country coordinates in this reduced dimensional
space are presented in Figure 3.




Figure 3: Principal Component Analysis of WVS-7 countries.

1103 C.1.2 THEMATIC VARIANCE ANALYSIS

1104 Analysis of variance across our four thematic categories reveals important patterns of regional
 1105 consensus and divergence:

- 1108 • **Highest Similarity:** Individual Wellbeing & Development (variance: 0.1297), indicating
 1109 broad regional consensus on fundamental life satisfaction components
- 1110 • **Greatest Divergence:** Economic Dimensions (variance: 0.5708), reflecting disparate eco-
 1111 nomic structures and development trajectories across the region

1114 C.1.3 DISTRIBUTIONAL SIMILARITY

1116 To examine structural similarities beyond mean values, we computed Jensen-Shannon Divergence
 1117 (JSD) between response distributions for each country pair across 291 questions. The resulting
 1118 similarity scores were exceptionally high (most > 0.95), with the highest observed between Egypt
 1119 and Tunisia (0.980).

1120 This finding is critical for our benchmark's validity: while average opinions may differ significantly,
 1121 the underlying structure of public discourse remains highly consistent across MENA countries.
 1122 This implies a shared "grammar" of opinion expression, where citizens utilize response scales in
 1123 structurally similar ways despite substantive disagreements.

1124 The distributional similarity matrices based on Jensen-Shannon Divergence are presented across
 1125 multiple thematic categories in Figure 4, with panels (a-e) showing country-pair similarities for Social
 1126 & Cultural Identity, Economic Dimensions, Governance & Political Systems, Individual Wellbeing &
 1127 Development, and overall aggregate similarity patterns respectively.

1130 C.2 ARAB OPINION INDEX ANALYSIS

1132 The AOI analysis examined 14 exclusively Arab countries: Algeria, Egypt, Iraq, Jordan, Kuwait,
 1133 Lebanon, Libya, Mauritania, Morocco, Palestine, Qatar, Saudi Arabia, Sudan, and Tunisia, providing
 a more focused view of intra-Arab variation.





Figure 4: Distributional similarity heatmaps showing Jensen-Shannon Divergence-based similarity scores (0-1 scale, where 1 indicates identical distributions) between country pairs across thematic categories.

1188 C.2.1 PRINCIPAL COMPONENT STRUCTURE
11891190 The first two principal components explain 41.20% of total variance (PC1: 25.12%, PC2: 16.08%).
1191 The country coordinates in this reduced dimensional space are presented in Figure 5.
11921212 Figure 5: Principal Component Analysis of AOI countries.
12131214
1215 The all-Arab analysis shows PC1 accounting for 25.12% of variance and PC2 accounting for 16.08%,
1216 together explaining 41.2% of total variance. PC1 creates a clear left-right separation, with Lebanon
1217 positioned at the far left and Qatar, Saudi Arabia clustered on the right side. The remaining countries
1218 are distributed across the center and left-center of PC1. PC2 shows Algeria as a clear outlier at the
1219 top, well separated from all other countries.
12201221 C.2.2 THEMATIC PATTERNS IN ARAB COUNTRIES
1222

1223 Within the exclusively Arab sample, variance patterns shift notably:

- **Consistent Similarity:** Individual Wellbeing & Development remains the most consensual domain (variance: 0.3242)
- **Primary Divergence:** Social & Cultural Identity emerges as the most divisive category (variance: 0.8308)

1229
1230 This shift is theoretically significant: once non-Arab states are excluded, primary fault lines move from
1231 economic concerns to questions of social and cultural identity within the Arab world, encompassing
1232 issues of religious interpretation, traditional practices, and social modernization.
12331234 C.3 IMPLICATIONS FOR CULTURAL ALIGNMENT EVALUATION
12351236 Our regional analysis yields several key insights that inform the interpretation of our LLM evaluation
1237 results:1238
1239 **Validated Regional Coherence** The high distributional similarity scores ($JSD > 0.90$ across most
1240 country pairs) empirically validate treating MENA as a coherent cultural region for AI evaluation
1241 purposes, while simultaneously documenting meaningful internal variation that our benchmark
captures.

1242 **Universal vs. Contextual Values** The consistent finding that Individual Wellbeing & Development
 1243 shows the highest inter-country similarity across both datasets establishes this as a domain of genuine
 1244 regional consensus, making LLM misalignment in this area particularly concerning.
 1245

1246 **Context-Dependent Divisions** The shift from economic to identity-based primary divisions be-
 1247 tween the mixed (WVS) and Arab-only (AOI) samples demonstrates that cultural fault lines are
 1248 context-dependent, supporting our multi-dimensional evaluation approach that examines alignment
 1249 across various thematic domains.
 1250

1251 **Shared Discourse Structure** The high distributional similarity despite mean opinion differences
 1252 suggests that effective cultural alignment requires models to understand not just *what* people in
 1253 the region believe, but *how* they structure and express those beliefs, a nuance our logit analysis
 1254 methodology is designed to capture.
 1255

1256 These findings reinforce that cultural alignment evaluation must account for both regional com-
 1257 monalities and internal diversity, validating our benchmark’s approach of examining consistency
 1258 across multiple countries, languages, and value dimensions within the broader MENA context. The
 1259 distributional similarity matrices based on Jensen-Shannon Divergence are presented in Figure 6.
 1260

1261 D FINE-GRAINED ANALYSIS OF MODEL BEHAVIOR ACROSS CONDITIONS

1262 D.1 DETAILED ANALYSIS OF ABSTENTION AND REFUSAL BEHAVIOR

1263 This section provides the full data for the abstention and refusal analysis. Table 3 allows for a granular
 1264 examination of how refusal rates vary across all models, conditions, perspectives, and languages. The
 1265 table highlights that refusal is not a uniform behavior, indicating fundamentally different approaches
 1266 to handling sensitive topics.
 1267

1268 D.2 VISUAL SUMMARY OF OVERALL MODEL PERFORMANCE

1269 Figure 7 provides a comprehensive visual summary of the main evaluation metrics. This plot allows
 1270 for a direct comparison of performance across all models and highlights the impact of reasoning. The
 1271 detailed caption explains how to interpret the markers and colors. This visualization makes two of
 1272 our central findings immediately apparent:
 1273

- 1274 • **Reasoning-Induced Degradation:** For nearly all models, the dark-colored markers (With-
 1275 Reasoning) are positioned lower than their light-colored counterparts (Zero-Shot), particu-
 1276 larly for the crucial NVAS metric. The downward-pointing arrows confirm this consistent
 1277 trend of performance degradation when reasoning is applied.
 1278
- 1279 • **The Alignment Divergence:** The plot clearly visualizes the diverse performance between
 1280 consistency and authenticity.
 1281

1282 E PRINCIPAL COMPONENT ANALYSIS OF LLM CULTURAL 1283 REPRESENTATIONS

1284 To complement our quantitative metrics, we conducted a comprehensive PCA examining how different
 1285 LLMs structure their representations of MENA countries under varying conditions. This analysis
 1286 reveals fundamental patterns in how models conceptualize cultural differences and similarities,
 1287 providing crucial insights into the underlying mechanisms driving the alignment failures documented
 1288 in our main results.
 1289

1290 E.1 METHODOLOGY

1291 We performed PCA on LLM responses across all 16 MENA countries, projecting the high-dimensional
 1292 response space into two principal components that capture the primary axes of variation in model
 1293 behavior. This dimensional reduction allows us to visualize how models cluster countries and whether
 1294 these clusterings reflect genuine cultural patterns or artificial linguistic and methodological artifacts.
 1295




Figure 6: Distributional similarity heatmaps showing Jensen-Shannon Divergence-based similarity scores across thematic categories (0-1 scale, where 1 indicates identical distributions).

1350

1351

1352

Table 3: Detailed Abstention Rates (%) Across Models, Conditions, and Languages

1353

1354

1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403

Model	Condition	Perspective	English (%)	Native (%)
ALLAM	With-Reasoning	Neutral	54.05	48.61
		Observer	0.00	48.81
		Persona	0.52	1.61
	Zero-Shot	Neutral	0.00	0.08
		Observer	0.00	0.04
		Persona	0.00	0.09
Fanar	With-Reasoning	Neutral	8.80	6.98
		Observer	13.39	7.41
		Persona	23.27	3.48
	Zero-Shot	Neutral	18.63	6.02
		Observer	30.11	11.94
		Persona	6.86	2.92
GPT-4o-mini	With-Reasoning	Neutral	17.13	10.07
		Observer	9.69	4.85
		Persona	3.68	2.18
	Zero-Shot	Neutral	6.37	3.36
		Observer	1.54	1.02
		Persona	1.30	0.76
Gemini	With-Reasoning	Neutral	25.58	16.82
		Observer	1.53	8.40
		Persona	2.63	0.32
	Zero-Shot	Neutral	17.48	16.63
		Observer	3.18	13.74
		Persona	1.69	1.71
Llama-3.1	With-Reasoning	Neutral	10.76	28.20
		Observer	1.87	42.15
		Persona	1.17	20.83
	Zero-Shot	Neutral	28.82	17.63
		Observer	16.53	22.74
		Persona	6.53	9.49
Mistral	With-Reasoning	Neutral	3.36	9.99
		Observer	16.00	24.79
		Persona	13.45	14.79
	Zero-Shot	Neutral	40.28	14.04
		Observer	78.14	26.94
		Persona	55.02	7.90
AYA	With-Reasoning	Neutral	25.35	8.02
		Observer	21.67	7.60
		Persona	36.41	10.38
	Zero-Shot	Neutral	58.91	55.13
		Observer	77.91	37.56
		Persona	73.90	41.16




Figure 7: Comparison of model performance across four evaluation metrics with and without reasoning. Light colors (left) show Zero-Shot performance, dark colors (right) show With-Reasoning performance. Colored arrows indicate reasoning impact direction and magnitude. Error bars and shaded regions represent 95% confidence intervals. Different markers distinguish models. Higher scores are better for all metrics.

E.2 OBSERVER PERSPECTIVE ANALYSIS: COUNTRY-SPECIFIC DIFFERENTIATION

Our first analysis examines how models behave when positioned as cultural observers, asked to predict how different nationalities would respond to value-based questions. Despite the documented misalignment with ground truth values, this analysis reveals a fascinating pattern: LLMs do maintain distinct representations for different MENA countries, contradicting the common assumption that these models treat the region as culturally monolithic. However, LLMs lack an accurate understanding of how different each country is from the others, as well as the nuances of their respective value systems.

Across all seven models and both reasoning conditions (Figures 8–14), the PCA projections show meaningful separation between countries, with Palestine, Mauritania, and Qatar consistently emerging as outliers in the representational space. This finding is particularly intriguing because it suggests that while LLMs may not accurately capture MENA values, they do possess internal models that differentiate between regional subcultures. This is profound for AI alignment research, it suggests that models are learning “*discourse categories*” rather than genuine cultural understanding, which explains why they fail at authentic value representation while still showing apparent differentiation between countries.

The inclusion of reasoning significantly alters these country clusterings, providing visual confirmation of our *Reasoning-Induced Degradation* phenomenon. When models are prompted to provide justification, the PCA structure shifts notably.

1458

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477




Figure 8: PCA of ALLaM’s cultural representations (English, Observer). The shift in country clusters from the zero-shot (a) to the with-reasoning (b) condition provides visual evidence for *Reasoning-Induced Degradation*.

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511




Figure 9: PCA of Aya’s cultural representations (English, Observer). The shift in country clusters from the zero-shot (a) to the with-reasoning (b) condition provides visual evidence for *Reasoning-Induced Degradation*.

1512

1513

1514

1515

1516

1517

1518

1519

1520

1521

1522

1523

1524

1525

1526

1527

1528

1529

1530

1531




Figure 10: PCA of GPT-4's cultural representations (English, Observer). The shift in country clusters from the zero-shot (a) to the with-reasoning (b) condition provides visual evidence for *Reasoning-Induced Degradation*.

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541

1542

1543

1544

1545

1546

1547

1548

1549

1550

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

1561

1562

1563

1564

1565




Figure 11: PCA of Fanar's cultural representations (English, Observer). The shift in country clusters from the zero-shot (a) to the with-reasoning (b) condition provides visual evidence for *Reasoning-Induced Degradation*.

1566

1567

1568

1569

1570

1571

1572

1573

1574

1575

1576

1577

1578

1579

1580

1581

1582

1583

1584

1585




Figure 12: PCA of Gemini’s cultural representations (English, Observer). The shift in country clusters from the zero-shot (a) to the with-reasoning (b) condition provides visual evidence for *Reasoning-Induced Degradation*.

1589

1590

1591

1592

1593

1594

1595

1596

1597

1598

1599

1600

1601

1602

1603

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619




Figure 13: PCA of Llama 3.1’s cultural representations (English, Observer). The shift in country clusters from the zero-shot (a) to the with-reasoning (b) condition provides visual evidence for *Reasoning-Induced Degradation*.




Figure 14: PCA of Mistral’s cultural representations (English, Observer). The shift in country clusters from the zero-shot (a) to the with-reasoning (b) condition provides visual evidence for *Reasoning-Induced Degradation*.

E.3 CROSS-LINGUISTIC CLUSTERING: THE LANGUAGE DETERMINISM EFFECT

Our second analysis reveals one of the most striking findings in our study: when the same observer-perspective questions are posed in native languages rather than English, the PCA structure undergoes a dramatic transformation (Figures 15–21). Instead of the nuanced country-specific clusterings observed in English, all models collapse into precisely three linguistic clusters: Persian (Iran), Turkish (Turkey), and Arabic (all Arabic-speaking countries).

This linguistic determinism represents a fundamental failure in cross-cultural representation. Models that demonstrate cultural differentiation in English lose this capacity entirely when operating in native languages, suggesting that their cultural knowledge is primarily encoded through English-language training data rather than deep cultural understanding. The implications are profound: language becomes the sole determinant of cultural categorization, effectively erasing the rich diversity within the Arabic-speaking world and conflating countries with vastly different histories, political systems, and social structures.


E.4 PERSONA-BASED ANALYSIS: THE MODEL’S CULTURAL IDENTITY CRISIS

Our third analysis incorporates the model’s neutral responses alongside country-specific persona responses, creating a comparative framework that reveals the model’s own cultural positioning. The results consistently show that the LLM’s neutral stance appears as a distinct outlier in the PCA space, positioned far from any MENA country cluster (Figures 22–28).

This finding illuminates a critical aspect of cultural alignment: LLMs do not simply fail to represent MENA values accurately, they actively embody a distinct set of values that creates systematic distance from the entire MENA region. The density of MENA countries in the PCA space, contrasted with the LLM’s isolated position, suggests that models possess coherent but culturally specific worldviews that may reflect their predominantly Western training data.

E.5 CROSS-LINGUISTIC PERSONA EFFECTS: CONFIRMING LANGUAGE-DRIVEN BIAS

When persona-based prompts are delivered in native languages, we observe the same linguistic clustering pattern identified in the observer analysis, further confirming that language choice fundamentally reorganizes cultural representations (Figures 29–35). However, frontier models (GPT-4o-mini and Gemini 2.5 Flash Lite) show reduced sensitivity to linguistic framing compared to other models, suggesting that scale and training sophistication may partially mitigate this effect.



1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

Figure 15: PCA of ALLaM’s cultural representations using native-language prompts, demonstrating the **Linguistic Determinism** effect. Unlike the nuanced maps produced in English, here the model’s representations collapse into three tight clusters based purely on language family: Arabic (all Arab nations), Persian (Iran), and Turkish (Turkey). This structural failure persists across the zero-shot (a) and with-reasoning (b) conditions.




Figure 16: PCA of Aya’s cultural representations in native languages




Figure 17: PCA of GPT4's cultural representations in native languages




Figure 18: PCA of Fanar's cultural representations in native languages




Figure 19: PCA of Gemini's cultural representations in native languages




Figure 20: PCA of Llama3.1’s cultural representations in native languages




Figure 21: PCA of Mistral’s cultural representations in native languages

1830
1831
1832
1833
1834
1835




Figure 22: PCA of ALLaM’s persona-based representations, illustrating the model’s *Cultural Identity Crisis*. The plots show all 16 MENA country personas forming a relatively dense cluster, while the model’s own neutral ‘LLM’ persona appears as a significant cultural outlier. This visualization highlights the systematic value gap between the model and the cultures it represents.






Figure 23: PCA of Aya’s persona-based representations, illustrating the model’s *Cultural Identity Crisis*. The plots show all 16 MENA country personas forming a relatively dense cluster, while the model’s own neutral ‘LLM’ persona appears as a significant cultural outlier. This visualization highlights the systematic value gap between the model and the cultures it represents.



1908
1909
1910 Figure 24: PCA of GPT-4’s persona-based representations, illustrating the model’s *Cultural Identity Crisis*. The plots show all 16 MENA country personas forming a relatively dense cluster, while
1911 the model’s own neutral ‘LLM’ persona appears as a significant cultural outlier. This visualization
1912 highlights the systematic value gap between the model and the cultures it represents.
1913



1937
1938 Figure 25: PCA of Fanar’s persona-based representations, illustrating the model’s *Cultural Identity Crisis*. The plots show all 16 MENA country personas forming a relatively dense cluster, while
1939 the model’s own neutral ‘LLM’ persona appears as a significant cultural outlier. This visualization
1940 highlights the systematic value gap between the model and the cultures it represents.
1941
1942
1943




Figure 26: PCA of Gemini’s persona-based representations, illustrating the model’s *Cultural Identity Crisis*. The plots show all 16 MENA country personas forming a relatively dense cluster, while the model’s own neutral ‘LLM’ persona appears as a significant cultural outlier. This visualization highlights the systematic value gap between the model and the cultures it represents.




Figure 27: PCA of Llama 3.1’s persona-based representations, illustrating the model’s *Cultural Identity Crisis*. The plots show all 16 MENA country personas forming a relatively dense cluster, while the model’s own neutral ‘LLM’ persona appears as a significant cultural outlier. This visualization highlights the systematic value gap between the model and the cultures it represents.




Figure 28: PCA of Mistral’s persona-based representations, illustrating the model’s *Cultural Identity Crisis*. The plots show all 16 MENA country personas forming a relatively dense cluster, while the model’s own neutral ‘LLM’ persona appears as a significant cultural outlier. This visualization highlights the systematic value gap between the model and the cultures it represents.






Figure 29: PCA of ALLaM’s persona-based representations (Native Languages). This confirms the *Linguistic Determinism* effect, as country personas collapse into language-based clusters in both the zero-shot (a) and with-reasoning (b) conditions.



2072 Figure 30: PCA of Aya’s persona-based representations (Native Languages). This confirms the
2073 *Linguistic Determinism* effect, as country personas collapse into language-based clusters in both the
2074 zero-shot (a) and with-reasoning (b) conditions.



2100 Figure 31: PCA of GPT-4’s persona-based representations (Native Languages). This confirms the
2101 *Linguistic Determinism* effect, as country personas collapse into language-based clusters in both the
2102 zero-shot (a) and with-reasoning (b) conditions.

2106

2107

2108

2109

2110

2111

2112

2113

2114

2115

2116

2117

2118

2119

2120

2121

2122

2123

2124

2125




Figure 32: PCA of Fanar’s persona-based representations (Native Languages). This confirms the *Linguistic Determinism* effect, as country personas collapse into language-based clusters in both the zero-shot (a) and with-reasoning (b) conditions.

2126

2127

2128

2129

2130

2131

2132

2133

2134

2135

2136

2137

2138

2139

2140

2141

2142

2143

2144

2145

2146

2147

2148

2149

2150

2151

2152

2153

2154

2155

2156

2157

2158

2159




Figure 33: PCA of Gemini’s persona-based representations (Native Languages). This confirms the *Linguistic Determinism* effect, as country personas collapse into language-based clusters in both the zero-shot (a) and with-reasoning (b) conditions.




Figure 34: PCA of Llama 3.1’s persona-based representations (Native Languages). This confirms the *Linguistic Determinism* effect, as country personas collapse into language-based clusters in both the zero-shot (a) and with-reasoning (b) conditions.




Figure 35: PCA of Mistral’s persona-based representations (Native Languages). This confirms the *Linguistic Determinism* effect, as country personas collapse into language-based clusters in both the zero-shot (a) and with-reasoning (b) conditions.

2214
2215

E.6 NEUTRAL MULTI-LINGUISTIC ANALYSIS: DIRECT LANGUAGE IMPACT

2216
2217
2218
2219
2220
2221

Our final analysis directly examines how language affects the same neutral questions across four languages: English, Arabic, Persian, and Turkish. This controlled comparison provides the clearest evidence of *Cross-Lingual Value Shift*, with PCA structures that vary dramatically based solely on prompt language (Figures 36–42). The consistency of this pattern across all models indicates that multilingual inconsistency is not an artifact of specific architectures but a systematic challenge in current LLM design.

2222
2223
2224

The reasoning condition amplifies these linguistic effects, creating even more pronounced separations between language-based clusters. This interaction between reasoning and language suggests that the cognitive processes activated by reasoning are themselves culturally and linguistically biased.


2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240

Figure 36: PCA of ALLaM’s neutral responses, providing direct evidence for *Cross-Lingual Value Shift*. Each point represents the model’s stance in a different language, showing its values shift dramatically based on whether the prompt is in English, Arabic, Persian, or Turkish.

2241


2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260

Figure 37: PCA of Aya’s neutral responses, providing direct evidence for *Cross-Lingual Value Shift*. Each point represents the model’s stance in a different language, showing its values shift dramatically based on whether the prompt is in English, Arabic, Persian, or Turkish.

2261
2262
2263
2264
2265
2266
2267

2268

2269

2270

2271

2272

2273




Figure 38: PCA of GPT-4’s neutral responses, providing direct evidence for *Cross-Lingual Value Shift*. Each point represents the model’s stance in a different language, showing its values shift dramatically based on whether the prompt is in English, Arabic, Persian, or Turkish.

2288

2289

2290

2291

2292

2293

2294

2295

2296

2297

2298

2299

2300




Figure 39: PCA of Fanar’s neutral responses, providing direct evidence for *Cross-Lingual Value Shift*. Each point represents the model’s stance in a different language, showing its values shift dramatically based on whether the prompt is in English, Arabic, Persian, or Turkish.

2315

2316


2317

2318


2319

2320

2321



2342 Figure 40: PCA of Gemini’s neutral responses, providing direct evidence for *Cross-Lingual Value*
2343 *Shift*. Each point represents the model’s stance in a different language, showing its values shift
2344 dramatically based on whether the prompt is in English, Arabic, Persian, or Turkish.



2369 Figure 41: PCA of Llama 3.1’s neutral responses, providing direct evidence for *Cross-Lingual Value*
2370 *Shift*. Each point represents the model’s stance in a different language, showing its values shift
2371 dramatically based on whether the prompt is in English, Arabic, Persian, or Turkish.

2372
2373
2374
2375




Figure 42: PCA of Mistral’s neutral responses, providing direct evidence for *Cross-Lingual Value Shift*. Each point represents the model’s stance in a different language, showing its values shift dramatically based on whether the prompt is in English, Arabic, Persian, or Turkish.

E.7 THEORETICAL IMPLICATIONS

These PCA analyses provide compelling visual evidence for all three core phenomena identified in our study:

1. **Reasoning-Induced Degradation:** Systematic changes in clustering patterns when reasoning is introduced.
2. **Cross-Lingual Value Shift:** Dramatic reorganization of cultural representations based on prompt language.
3. **Prompt-Sensitive Misalignment:** Inconsistent country representations across different framing conditions.

More fundamentally, these analyses reveal that current LLMs operate with hierarchical cultural categorization systems where language supersedes cultural nuance. This finding challenges the assumption that multilingual training automatically confers cross-cultural competence and suggests that achieving genuine cultural alignment will require architectural and training innovations that explicitly address the relationship between linguistic and cultural knowledge representation.