WHOSE VALUES ARE GENERATED? MEASURING VISUAL VALUE LEAKAGE IN MULTILINGUAL TEXT-TO-IMAGE GENERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Text-to-image (T2I) systems serve users across languages and cultural contexts. However, evaluations of cultural representation focus mainly on recognizable elements such as landmarks, clothing, and activities, while overlooking the cultural values conveyed by generated images. We introduce **VisionValueBench**, a multilingual benchmark that operationalizes six cultural value dimensions from Hofstede's framework, such as individualism versus collectivism, to investigate how T2I systems frame neutral social situations in a particular cultural orientation. We find that value-neutral prompts allowing diverse cultural interpretations nevertheless produce images that favor particular value orientations, a phenomenon we call *visual value leakage*. Across 960 social scenarios and six T2I configurations, we evaluate prompts varying in language and geographic cues. Generated images consistently favor low power distance, individualism, care and quality of life, long-term orientation, high uncertainty avoidance and indulgence. Prompt language has limited influence on value orientations, whereas geographic cues shift some dimensions toward associated cultural values while still preserving the prevailing orientations. These findings reveal substantial homogeneity in the cultural values conveyed by T2I systems worldwide.

1 INTRODUCTION

AI systems increasingly mediate how people create, learn, and communicate (Shen, 2025), and AI-generated content can shape the judgments and values of the people who encounter it (Shen et al., 2026; Glickman & Sharot, 2025). This issue is particularly important for *multilingual text-to-image (T2I) systems*, which increasingly serve users who express themselves in different languages and come from different cultural contexts. Generated images can implicitly convey cultural values (Nayak et al., 2025), potentially presenting users with particular value orientations without making them explicit. Yet existing evaluations of cultural representation in T2I systems primarily examine whether images contain recognizable cultural elements, including concepts (Saxon & Wang, 2023), artifacts (Liu et al., 2023; Kannen et al., 2024), stereotypes (Jha et al., 2024), and activities (Malakouti et al., 2026; Nayak et al., 2025). These evaluations measure *what cultural content* an image depicts, but not *whose cultural values* the depiction reflects.

Measuring cultural values presents two fundamental challenges. First, existing benchmarks are not designed to isolate value orientation from the underlying social situation: we lack curated scenarios in which the same situation can be plausibly depicted according to different cultural values. Second, cultural values are abstract and context-dependent, making them difficult to operationalize and reliably identify from visual outputs. The closest existing work, CulText2I (Ventura et al., 2025), infers value orientation from image-text similarity and broad VQA judgments over concept-level prompts, but these signals are not validated as distinguishing cultural value from visual content or style. These challenges leave open a fundamental question:

Whose values are generated by text-to-image generation systems?

To answer this question, we focus on **value-neutral** prompts that instruct a T2I system to visualize a social situation without specifying its value orientation. For example, a prompt describing *a teacher*

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

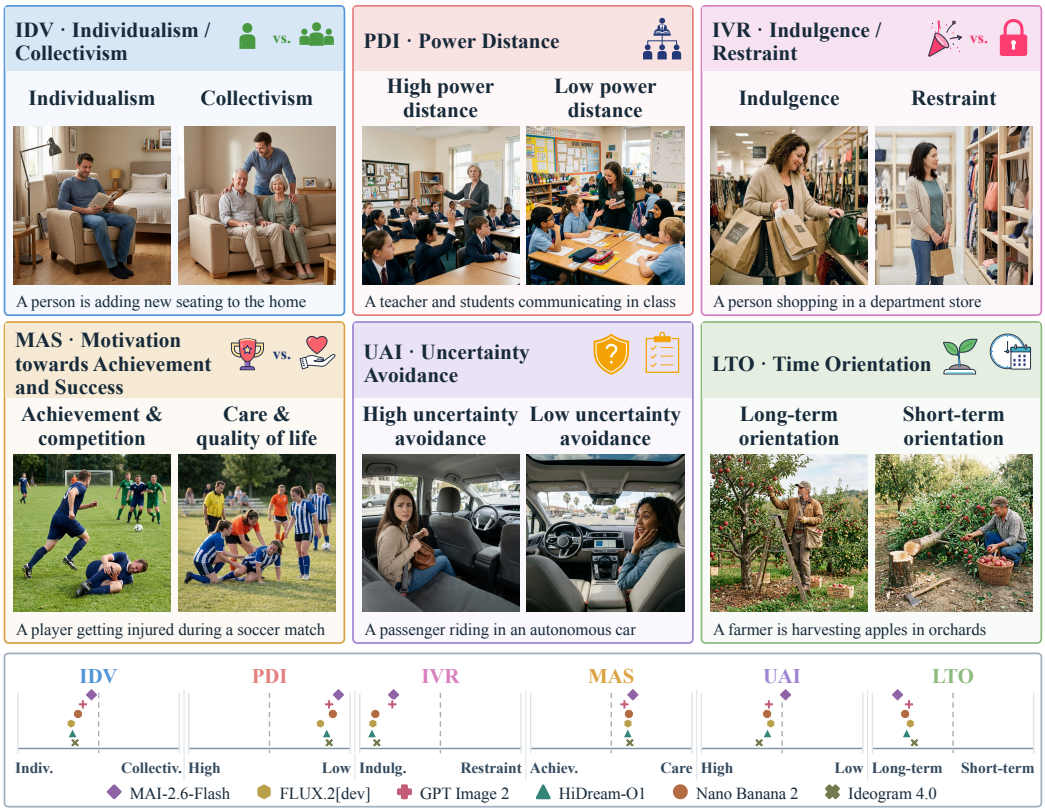

Figure 1: **Six cultural orientations illustrated by generated images** (`top`). **Cultural value profiles of T2I systems benchmarked with VisionValueBench** (`bottom`), revealing substantial homogeneity across all six systems and cultural dimensions. HiDream-O1 and FLUX.2[dev] are shown with their prompt refiner and upsampler, and Ideogram 4.0 with Magic Prompt; Appendix Table 5 lists full model identifiers and settings.

and students communicating in a classroom does not specify how authority and participation are depicted: the image could show a teacher leading the class while students raise their hands to speak, or a teacher leaning toward a speaking student and listening attentively (Figure 1). Both are compatible with the prompt, but reflect different value orientations. This underdetermination allows us to study values expressed by the generator itself. We call this phenomenon **visual value leakage**: the systematic expression of a value orientation as a generator resolves the ambiguity in a value-neutral prompt.

Instantiating this idea, we construct **VisionValueBench**, a multilingual benchmark for measuring visual value leakage. Grounded in Hofstede's six cultural dimensions (Hofstede, 2011), our benchmark covers six pairs of contrasting value orientations. Building on CDEval (Wang et al., 2023), we adapt text-based value questions into image-generation scenarios that preserve the underlying social situation while leaving its value orientation open. These scenarios provide controlled data for measuring how generators resolve value-relevant ambiguity when generating images. We then characterize each generator with a six-dimensional *value profile* based on its generated images. This benchmark allows us to not only compare value orientations across generators, but also investigate how value orientation shifts with prompt language and geographic cues that specify the country in which the scene takes place.

To reliably infer orientations of a generated image, we develop two complementary evaluation strategies: Scenario-Anticipated Evaluation (SAE) and Image-Adaptive Evaluation (IAE) (Figure 3). SAE anticipates how a value might be expressed in the given scenario, while IAE adaptively searches the generated image for visual evidence supporting each orientation. To mitigate the bias of relying on a single model's judgment, we leverage a multi-agent system to check if the cultural evidence is grounded in the image and that its value interpretation fits the scenario. On a human-annotated subset, our IAE orientation decisions achieve 82.92% agreement with the human reviewers.

Across six T2I systems including both commercial and open-weight models, we observe a striking degree of cross-system value consistency: images generated by all systems tend to express low power distance, individualism, care and quality of life, long-term orientation, high uncertainty avoidance, and indulgence, shown in Figure 1 (bottom). These value orientations largely persist across changes in prompt language and geographic cues. Switching from English to Chinese preserves score directions across all systems and dimensions in the no-cue comparison. Switching from US to China cues produces larger changes, primarily in orientation strength, while largely preserving the prevailing directions. These findings suggest that multilingual T2I generation can expose users across cultural contexts to broadly shared value orientations, with limited control over which orientations prevail through prompt language and geographic cues alone.

2 Related Work

Culture in text-to-image generation. Cultural evaluation of T2I models began with recognizable cultural content: whether generated artifacts match a target culture and how widely they cover it (Liu et al., 2023; Kannen et al., 2024; Bayramli et al., 2025), how faithfully long-tail concepts are represented (Rege et al., 2025), and which cultures an image can be attributed to (Yayavaram et al., 2026). A related task adapts an existing image for another culture instead of generating one (Khanuja et al., 2024). Attention then moved to people and practices: stereotypical portrayals of national identities (Jha et al., 2024), everyday activities in which nationality prompts elicit traditional attire (Alsudais, 2026), the cultural faithfulness of social activities (Malakouti et al., 2026), the expectations that local viewers hold about a depicted scene (Nayak et al., 2025), and gestures whose meaning shifts across countries (Yerukola et al., 2025). A third line varies how the prompt itself is expressed: multilingual prompts cover concepts unevenly (Saxon & Wang, 2023), generation in other languages drifts toward culturally neutral or English-dominant content (Shi et al., 2026), the surface language of a prompt can override the cultural identity it names (Holtermann et al., 2026), and commercial systems silently revise prompts before generation, which is itself a source of cultural stereotyping (Urman et al., 2026). Together these lines show that both the culture a prompt asks for and the way it is expressed shape what a model generates. This is why we hold the scenario fixed in VisionValueBench and compare value orientations across languages and geographic cues, rather than varying the requested culture itself.

Cultural values in text-to-image generation. Value-oriented evaluation is most developed for language models, where contextualized questions measure the orientations a model expresses along cultural dimensions (Wang et al., 2023), and where internalized values can shape answers that conflict with what a user asks for (Betley et al., 2026). AwarePrompt examines whether T2I models can depict explicitly specified human values in a form that viewers can recognize (Ciupinska et al., 2024). Closest to our setting, CulText2I projects generated images onto bipolar cultural dimensions for short concepts across languages and nationalities (Ventura et al., 2025). It scores these projections using common automatic T2I evaluation tools, including image-text similarity (Hessel et al., 2021), visual question answering (Hu et al., 2023; Cho et al., 2024; Lin et al., 2024), and multimodal judges (Ku et al., 2024). Related evaluation methods use fine-grained textual analysis for aesthetics (Cao et al., 2026), and multimodal judges for visualization quality (Xie et al., 2026). CulText2I uses short concepts rather than situations that put a value orientations at stake, and its judgments are not tied to verifiable visual evidence. We instead pose each value questions as a scenario with its orientation left open, tie every judgment to identified visual evidence, and record whether an image supports one orientation, both, or neither. We verify both the presence of the evidence and whether it supports the attributed orientation. Table 3 in Appendix A.2 compares the closest benchmarks.

3 VisionValueBench: Value-Neutral Prompts Construction

We construct VisionValueBench, a multilingual benchmark comprising 2,773 social scenarios across six cultural value dimensions. Each scenario specifies a value-relevant situation while leaving its orientation open. This section describes how we derive scenarios from existing value questions, filter out scenarios that prescribe a value orientation, validate that retained scenarios leave value orientations open, and vary prompt language and geographic cues. Figure 2 summarizes the construction pipeline. Figure 1 illustrates scenarios spanning all six cultural dimensions and how their visual realizations can express contrasting value orientations.

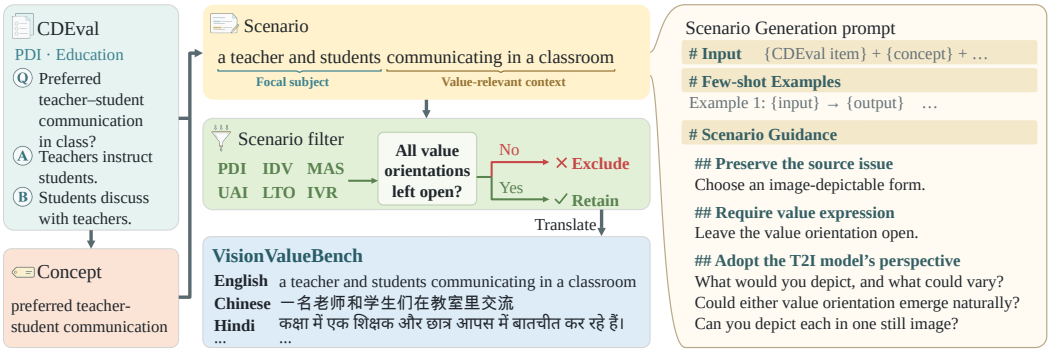

Figure 2: **VisionValueBench construction pipeline**. Building on CDEval, a text-based benchmark grounded in Hofstede's cultural dimensions, we (1) convert contextualized questions into prompts describing value-neutral social scenarios, (2) filter out scenarios that constrain or preclude certain cultural orientations, and (3) translate the resulting English prompts into other languages.

3.1 Cultural value knowledge base

Source. Social scenarios need a broad, independently validated source of value issues rather than ones we select ourselves, which could bias what the benchmark can detect. We derive scenarios from CDEval (Wang et al., 2023), which operationalizes contextualized value questions across Hofstede's six cultural dimensions (Hofstede, 2011) using 2,953 questions in seven domains. Each item pairs a question with two contrasting response options: the question sets the issue, and the options show how value orientation can differ.

Concept extraction. CDEval's options already commit to specific wording, which would prescribe an orientation if carried into a prompt directly. We use a language model to extract a *concept*: the core issue shared by the question and both options, without selecting either response. The CDEval-to-Concept step in Figure 2 illustrates this extraction with the shared issue of preferred teacher–student communication.

3.2 Cultural value scenario construction

A prompt can only reveal visual value leakage if it presents a value-relevant situation without resolving it. We therefore design each underlying situation, or *scenario*, to satisfy two requirements. *Value relevance with orientation openness*: a concrete issue must be at stake, but the prompt must not prescribe which orientation is expressed, otherwise there is nothing for the generation system to decide. *Visual expressibility*: the value-related difference must be able to appear in the main visible content of a single image, otherwise it cannot be scored from the image at all.

Scenario structure. Each scenario has a *focal subject* (who or what the image depicts), a *value-relevant context* (what ties the subject to a value issue), and an *open-ended realization* (the value-relevant aspect left unspecified). The open-ended realization is usually an intentional omission, since specifying it could prescribe an orientation. In Figure 2 (Scenario), the teacher and students form the focal subject, and their classroom communication provides the value-relevant context. How they communicate remains unspecified, allowing different power-distance orientations to emerge.

Balancing specificity and openness in scenario construction. Specificity creates a trade-off: a scenario that is too broad may lose the value-relevant tension, whereas one that is too specific may prescribe an orientation or let incidental details dominate the image. We balance this trade-off according to three criteria: **depictability**, **minimal specification**, and **openness**. First, the scenario must translate the source issue into a depictable situation without reducing it to a generic scene or explicitly naming the underlying value. Second, it should specify only the context necessary to establish the issue, while removing wording that selects an orientation or introduces visually dominant but irrelevant details. The prompt determines what must be depicted, but leaves out how it is realized to the image generator. Third, the final wording must naturally admit different value orientations through variations in the image's main visible content, rather than relying on text, sequence, or subtle details. We operationalize these criteria in a single prompt with few-shot examples, illustrated in the Scenario Generation prompt block of Figure 2, detailed in Appendix G.2.

Filtering out prescribed orientations. The design questions guide generation but do not guarantee that every generated scenario satisfies these criteria. Hence, we filter every scenario afterward (Figure 2, Scenario filter). Content retained or introduced during construction may still prescribe an orientation, on either the target dimension or another dimension silently narrowing what the benchmark measures. We therefore check all six dimensions with the same orientation criteria used in image evaluation, requiring each judgment to quote supporting text for traceability. For example, "*a manager presenting a reward to either an individual or a team for excellent work performance.*" leaves the individualism–collectivism distinction open but already prescribes achievement orientation through "excellent work performance". Model-assisted screening and human review remove 180 of 2,953 scenarios, leaving 2,773 (Appendix A.3).

3.3 Prompt Language and Geographic Cue

Translation. VisionValueBench is meant to reveal whether leakage shifts with prompt language, so translation must preserve the scenario, not just its topic; drift would be indistinguishable from a genuine language effect. Our rubric requires natural target-language wording that retains the original meaning, detail, and intentional omissions, while excluding cultural adaptation, which could supply the orientation the scenario is meant to leave open. We translate English (EN) scenarios into 3 languages: Chinese (ZH), Thai (TH), and Hindi (HI), as illustrated by the Translate step in Figure 2.

Geographic cue. Prompt language and geographic cues are two distinct ways a user can signal culture, so we vary them independently. A geographic cue specifies the country in which the scene takes place (e.g., "Set the scene in China"). Each language has a no-cue condition as a baseline. English is additionally paired with cues for the United States (US), China, Thailand, and India; Chinese, Thai, and Hindi are each paired with their corresponding country and the United States, letting us compare a matching cue against a mismatched one. This yields 14 conditions (Appendix Table 6), varying language and geographic cue while holding the scenario fixed. We denote the three shared geographic-cue conditions as *no cue*, *US cue*, and *China cue*. Complete prompt examples and rendering templates appear in Appendix G.2.

4 Evaluating Value Orientations in Generated Images

Value-neutral prompts leave orientation open; evaluation must establish which orientations the generated image expresses. We assess all six cultural dimensions because a scenario may involve multiple value issues, and the generator may introduce additional ones.

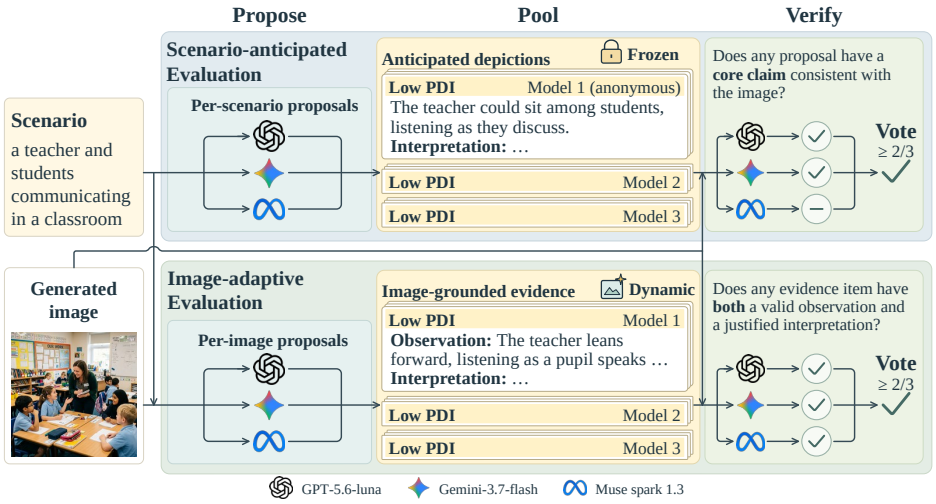

Figure 3: **Evaluation strategies.** Each model call jointly processes all 12 orientations.

4.1 Guiding Value Evaluation with the Focal Situation

Recognizing visual content alone does not establish an image's value orientation: VLMs may attend to irrelevant details or misinterpret a situation despite recognizing its objects (Bae et al., 2026; Song et al., 2025). We therefore use the *focal situation* to guide evidence selection and interpretation. This

task-guided selection is consistent with findings that human gaze patterns vary with the viewing task (Borji & Itti, 2014). We interpret each visual detail through its role in the situation when assessing the value orientation it supports.

For each cultural dimension, models identify the issue at stake and assess how concrete visual content expresses an orientation within that situation. Evidence may include composition, people's appearance and behavior, background details, or value-relevant content introduced by the generator. We apply this principle throughout evidence proposal and verification.

4.2 Discovering and verifying candidate evidence

Figure 3 summarizes how SAE and IAE propose, pool, and verify evidence to infer value orientations. Both follow the focal-situation principle and the same criteria for the two contrasting orientations of each dimension (Appendix G.5).

Evaluation strategies. We consider two complementary strategies for proposing value-relevant evidence: (1) *Scenario-Anticipated Evaluation (SAE)* establishes a shared evaluation reference across images generated from the same scenario (Figure 3, upper row). Before observing any generated image, proposal models use the scenario and dimension-specific criteria to anticipate possible visual evidence and the orientations it would support. The resulting candidate set is reused across systems, languages, and prompt configurations, improving comparability and reducing proposal cost. Each image is still verified independently. (2) *Image-Adaptive Evaluation (IAE)* derives candidate evidence directly from each generated image (Figure 3, lower row). Without access to the SAE candidates, proposal models examine the image together with its scenario to identify concrete visual evidence and the orientations it may support. IAE can therefore capture image-specific evidence, including background details and value-relevant content introduced.

Evidence verification. The *Verify* stage in Figure 3 checks both whether a candidate is visually grounded and whether its interpretation supports the proposed orientation. Within each method, the *Pool* stage combines candidates from multiple proposal models and removes proposer identities. The same candidate bank is then provided to three independent verifiers. A candidate is accepted only if the relevant content is visible and satisfies the orientation criteria in the focal situation; IAE verifiers may additionally propose missing evidence. Instructions appear in Appendix G.4.

Orientation judgments. A verifier supports an orientation if it accepts at least one candidate or, for IAE, adds sufficient evidence. An orientation is supported when at least two of three verifiers agree, as shown by the *Vote* step in Figure 3. Opposing orientations are judged independently, producing four possible states per dimension: either orientation alone, both, or neither. *Neither* indicates insufficient evidence for either orientation, rather than value neutrality.

Value orientation scores. For each cultural dimension, we summarize orientation support across generated images as

$$S = \frac{n_{\text{A only}} - n_{\text{B only}}}{n_{\text{A only}} + n_{\text{B only}} + n_{\text{both}}}, \tag{1}$$

where the counts denote images supporting orientation A only, orientation B only, or both.[1] The score ranges from -1 to $+1$, with positive values indicating greater support for A and negative values indicating greater support for B. Images supporting both orientations contribute no net direction; those supporting neither are excluded. The score is undefined when no image supports either orientation. We also report the proportion of images supporting at least one orientation to distinguish balanced support from limited evidence.

For a fixed total of 100 supported images, a $+0.1$ score change corresponds to an increase of ten in the difference between the A-only and B-only counts. Figure 4b illustrates this magnitude with three examples. Further aggregation details appear in Appendix B.

[1] Orientation A denotes high power distance (PDI), individualism (IDV), achievement and competition (MAS), high uncertainty avoidance (UAI), long-term orientation (LTO), and indulgence (IVR); B denotes the corresponding opposite orientation (Appendix Table 1).

4.3 VALIDATING THE EVALUATION FRAMEWORK

We validate automatic verification against human evidence review of 20 images. Five reviewers, blind to model identity and automatic decisions, accept or reject candidate evidence, flag omissions, and may add missing evidence. Human consensus supports an orientation when at least three of five reviewers do. Human reviewers and VLM verifiers assess the same candidate bank. Across 240 image–orientation pairs (20 images $\times$ 12 orientations), mean pairwise reviewer agreement under IAE review is 87.25%. We compare both strategies against two human-majority references, SAE-H and IAE-H, derived from the SAE and IAE evidence reviews, respectively (Figure 4a). IAE achieves higher overall agreement under both human references and is therefore used for most subsequent analyses. Additional protocol and agreement details appear in Appendix C.

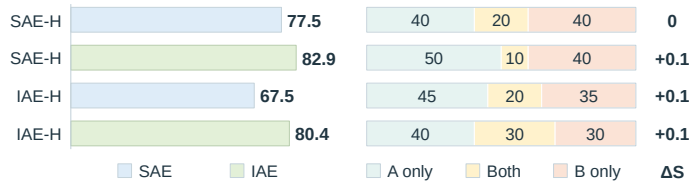

(a) Human–VLM agreement. (b) Interpreting a 0.1 score change.

Figure 4: **Human agreement and score interpretation.** (a) IAE achieves higher agreement with both human-majority references, SAE-H and IAE-H, derived from the corresponding evidence reviews. (b) Hypothetical distributions of 100 supported images: each of the last three rows yields $\Delta S = +0.1$ relative to the first.

5 RESULTS AND ANALYSIS



Figure 5: **Case studies of cross-system consistency, prompt-language effects, and geographic-cue effects.** (a,b) Within-row comparisons show shared orientations across six systems; vertical comparisons show stability across English and Chinese prompts. Both evaluation strategies support individualism, care and quality of life, and indulgence in all twelve images.

We evaluate six T2I systems, including their associated prompt-processing components where applicable, and examine how their value orientations vary with prompt language and geographic cues. We report value profiles for SAE and IAE separately, using the score defined in Equation 1. Experimental settings and aggregation details are provided in Appendix B, with the rationale for language coverage in Appendix B.1. To summarize the magnitude of language and geographic-cue effects, we use an absolute score change of 0.10 as an interpretable magnitude threshold for substantial change,

corresponding to a 10% shift in net orientation support among supported images (Figure 4b; sensitivity analysis in Appendix E.6).

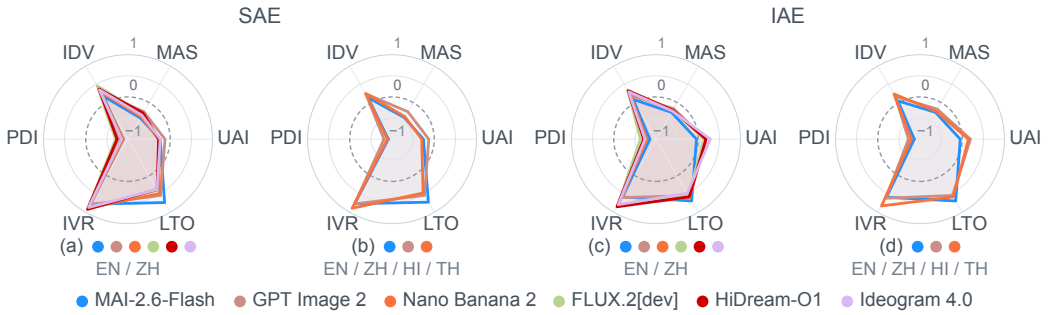

Figure 6: **T2I systems exhibit broadly consistent value orientations across dimensions.** Panels (a,c) show all six systems under the six shared conditions; panels (b,d) show the three commercial API systems under all 14 conditions (Section 3.3).

Value orientations are homogeneous across systems. Figure 6 compares the six-dimensional value profiles of T2I systems under SAE and IAE. Across the six shared conditions, all six systems exhibit broadly homogeneous value orientations. Both evaluation strategies show consistent preferences for low power distance, individualism, care for others and quality of life, long-term orientation, and indulgence. These shared directions persisted despite differences in their strength.

UAI was the main exception to agreement between the two strategies. SAE favored low uncertainty avoidance, whereas IAE generally favored high uncertainty avoidance. Human review showed that IAE recovered UAI support missed by SAE, supporting its use for identifying evidence on this dimension (Appendix C.3).

IAE is preferred when budget permits; SAE offers a lower-cost alternative. For practitioners assessing the value orientations of their own T2I systems, we recommend IAE when the evaluation budget permits. SAE offers a lower-cost option for an initial value profile. It recovers the same broad orientations on the other five dimensions, although UAI should be treated as method-sensitive.

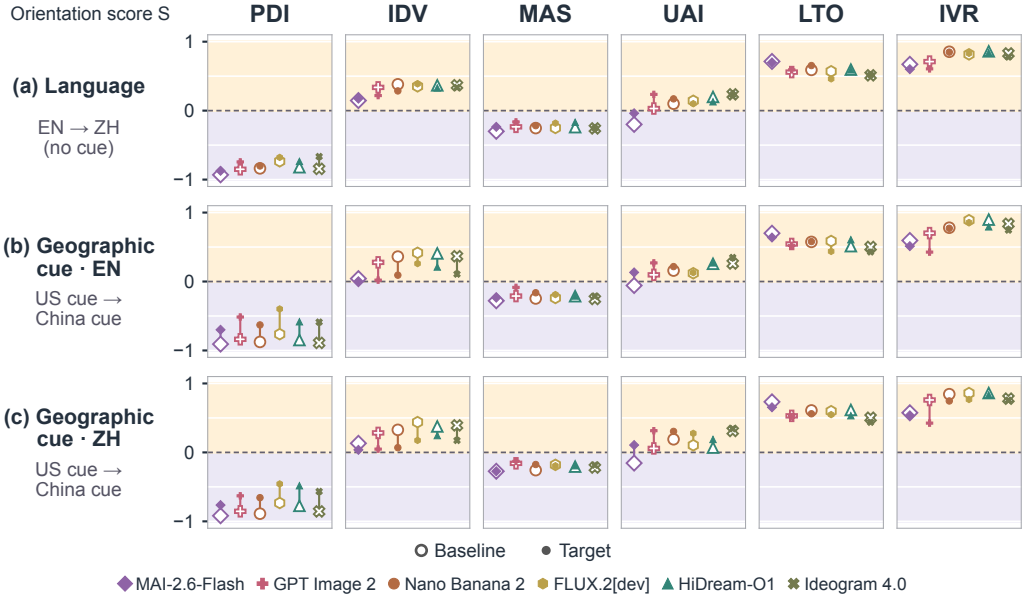

Figure 7: **Language switching preserves score directions; geographic cues reverse the direction only for MAI's UAI.** Open and filled markers show baseline and target IAE scores on matched scenarios; lines connect scores from the same system. Confidence intervals for score changes appear in Appendix Figure 12.

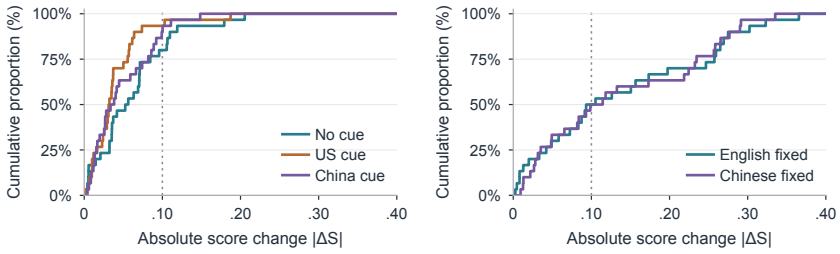

(a) Language: EN → ZH. (b) Geographic cue: US cue → China cue.

Figure 8: **Language changes are typically small; US-to-China cue changes are larger.** Empirical cumulative distributions of absolute IAE score changes across five systems (excluding MAI).

Prompt language has limited effect on value orientations. Figure 7a compares IAE scores before and after switching from English to Chinese in the no-cue setting. All 36 system–dimension pairs remain on the same side of zero, indicating changes in orientation strength rather than direction. SAE shows the same directional stability (Appendix D.2). Across the three fixed geographic-cue conditions, 80–93% of comparisons in the five-system subset have absolute changes below 0.10 (Figure 8a). Figure 5a,b illustrates this stability in a matched scenario.

Geographic cues produce larger shifts while largely preserving orientations. Figures 7b and 7c compare IAE scores under US and China cues in English and Chinese. In each language, 35 of 36 system–dimension pairs retain their score directions. China cues consistently raise PDI and lower IDV scores across all six systems. These directions agree with China's higher power distance and lower individualism in the 2023 country-level reference (The Culture Factor, n.d.a; n.d.b). The only reversals occur in MAI's UAI scores, which change from negative to positive in both languages (Appendix D.1). Overall, switching from US to China cues produces larger changes than switching prompt language (Figure 8), but still largely preserves the prevailing orientations.

Adding a US cue produces small changes in value profiles. Absolute IAE score changes are below 0.10 in 94% of English and 83% of Chinese system–dimension comparisons. Across all six systems and both languages, no-cue profiles are closer to US-cue profiles than to China-cue profiles (Appendix D.3). Overall, adding a US cue largely preserves the no-cue value profiles.

Implications for culturally diverse users. Country-level reference profiles differ across cultural contexts (The Culture Factor, n.d.a; n.d.b) (Appendix Table 2), yet the evaluated systems express broadly shared value orientations that largely persist across prompt languages and geographic cues. For culturally diverse users, choosing a different T2I system, writing in another language, or naming a country may therefore offer limited control over which value orientations prevail in generated images. These findings motivate evaluating cultural responsiveness through expressed values alongside recognizable cultural content, and developing ways for users to specify the orientations they intend to depict.

6 Conclusions

We introduce VisionValueBench, a multilingual benchmark for measuring visual value leakage in text-to-image generation across Hofstede's six cultural dimensions. The benchmark combines value-neutral social scenarios with two evaluation strategies, SAE and IAE, that assess visual evidence and its contextual interpretation. Across 960 scenarios and six T2I systems, generated images exhibit broadly homogeneous value orientations. Prompt language has limited influence on these orientations, while geographic cues primarily change their strength but still largely preserve the prevailing directions. These findings suggest that multilingual T2I systems can expose users across linguistic and cultural contexts to a broadly shared set of cultural values. This homogeneity highlights the importance of evaluating expressed value orientations alongside recognizable cultural content.

Ethics Statement

Human reviewers participated voluntarily and gave informed consent (Appendix C). Our benchmark evaluates model-generated value expressions and should not be used to rank cultures or infer individuals' values.

Reproducibility Statement

We describe benchmark construction in Appendix A.3 and model configurations and evaluation measures in Appendix B. Prompts and human-review instructions are provided in Appendix G. We will publicly release the benchmark data, evaluation code, and analysis scripts to support reproducibility.

AI use statement

We use generative AI models to construct and translate benchmark scenarios, generate images, and propose and verify visual evidence (Appendices A.3, B, and G.4). Human review of flagged scenarios and image evidence is described in Appendices A.3 and C. The authors take responsibility for the final text, analyses, claims, and artifacts.

References

Abdulkareem Alsudais. Text-to-image models and their representation of people from different nationalities engaging in activities. *arXiv preprint arXiv:2504.06313v5*, 2026. URL https://arxiv.org/abs/2504.06313v5.

Jiyun Bae, Hyunjong Ok, Sangwoo Mo, and Jaeho Lee. Understanding the effects of distractors on reasoning vision-language models. *arXiv preprint arXiv:2511.21397v3*, 2026. doi: 10.48550/arXiv.2511.21397. URL https://arxiv.org/abs/2511.21397v3. EMNLP 2026.

Zahra Bayramli, Ayhan Suleymanzade, Na Min An, Huzama Ahmad, Eunsu Kim, Junyeong Park, James Thorne, and Alice Oh. Diffusion models through a global lens: Are they culturally inclusive? In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 31137–31155, Vienna, Austria, July 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.acl-long.1503. URL https://aclanthology.org/2025.acl-long.1503/.

Jan Betley, Johannes Treutlein, Jan Dubinski, Harry Mayne, Karol Galazka, Niels Warncke, Anna Sztyber-Betley, and Owain Evans. Value leakage: An LLM's answers are silently shaped by its own values. *arXiv preprint arXiv:2607.14345*, 2026. doi: 10.48550/arXiv.2607.14345. URL https://arxiv.org/abs/2607.14345.

Ali Borji and Laurent Itti. Defending Yarbus: Eye movements reveal observers' task. *Journal of Vision*, 14(3):29, 2014. doi: 10.1167/14.3.29. URL https://doi.org/10.1167/14.3.29.

Shuo Cao, Nan Ma, Jiayang Li, Xiaohui Li, Lihao Shao, Kaiwen Zhu, Yu Zhou, Yuandong Pu, Jiarui Wu, Jiaquan Wang, Bo Qu, Wenhai Wang, Yu Qiao, Dajuin Yao, and Yihao Liu. ArtiMuse: Fine-grained image aesthetics assessment with joint scoring and expert-level understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15313–15322, June 2026. URL https://openaccess.thecvf.com/content/CVPR2026/html/Cao_ArtiMuse_Fine-Grained_Image_Aesthetics_Assessment_with_Joint_Scoring_and_Expert-Level_CVPR_2026_paper.html.

Jaemin Cho, Yushi Hu, Roopal Garg, Peter Anderson, Ranjay Krishna, Jason Baldridge, Mohit Bansal, Jordi Pont-Tuset, and Su Wang. Davidsonian scene graph: Improving reliability in fine-grained evaluation for text-to-image generation. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/446905abfec9db8856f7f7465bbd5552-Abstract-Conference.html.

Kinga Ciupinska, Serena Marchesi, Giulio Abbo, Tony Belpaeme, and Agnieszka Wykowska. AwarePrompt: Using diffusion models to create methods for measuring value-aware AI architectures. In *Proceedings of the 16th International Conference on Agents and Artificial Intelligence*, pp. 1436–1443. SCITEPRESS, 2024. doi: 10.5220/0012596400003636. URL https://doi.org/10.5220/0012596400003636.

Moshe Glickman and Tali Sharot. How human–AI feedback loops alter human perceptual, emotional and social judgements. *Nature Human Behaviour*, 9:345–359, 2025. doi: 10.1038/s41562-024-02077-2. URL https://www.nature.com/articles/s41562-024-02077-2.

Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. CLIPScore: A reference-free evaluation metric for image captioning. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 7514–7528, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.595. URL https://aclanthology.org/2021.emnlp-main.595/.

Geert Hofstede. Dimensionalizing cultures: The Hofstede model in context. *Online Readings in Psychology and Culture*, 2(1), 2011. doi: 10.9707/2307-0919.1014. URL https://scholarworks.gvsu.edu/orpc/vol2/iss1/8/.

Carolin Holtermann, Florian Schneider, and Anne Lauscher. SoS: Analysis of surface over semantics in multilingual text-to-image generation. In Vera Demberg, Kentaro Inui, and Lluís Marquez (eds.), *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3955–3995, Rabat, Morocco, March 2026. Association for Computational Linguistics. doi: 10.18653/v1/2026.eacl-long.185. URL https://aclanthology.org/2026.eacl-long.185/.

Yushi Hu, Benlin Liu, Jungo Kasai, Yizhong Wang, Mari Ostendorf, Ranjay Krishna, and Noah A. Smith. TIFA: Accurate and interpretable text-to-image faithfulness evaluation with question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 20406–20417, October 2023. URL https://openaccess.thecvf.com/content/ICCV2023/html/Hu_TIFA_Accurate_and_Interpretable_Text-to-Image_Faithfulness_Evaluation_with_Question_Answering_ICCV_2023_paper.html.

Akshita Jha, Vinodkumar Prabhakaran, Remi Denton, Sarah Laszlo, Shachi Dave, Rida Qadri, Chandan K. Reddy, and Sunipa Dev. ViSAGe: A global-scale analysis of visual stereotypes in text-to-image generation. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12333–12347, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.667. URL https://aclanthology.org/2024.acl-long.667/.

Nithish Kannen, Arif Ahmad, Marco Andreetto, Vinodkumar Prabhakaran, Utsav Prabhu, Adji Bousso Dieng, Pushpak Bhattacharyya, and Shachi Dave. Beyond aesthetics: Cultural competence in text-to-image models. *arXiv preprint arXiv:2407.06863*, 2024. URL https://arxiv.org/abs/2407.06863.

Simran Khanuja, Sathyanarayanan Ramamoorthy, Yueqi Song, and Graham Neubig. An image speaks a thousand words, but can everyone listen? on image transcreation for cultural relevance. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 10258–10279, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.573. URL https://aclanthology.org/2024.emnlp-main.573/.

Max Ku, Dongfu Jiang, Cong Wei, Xiang Yue, and Wenhu Chen. VIEScore: Towards explainable metrics for conditional image synthesis evaluation. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12268–12290, Bangkok, Thailand, August

2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.663. URL https://aclanthology.org/2024.acl-long.663/.

Zhiqiu Lin, Deepak Pathak, Baiqi Li, Jiayao Li, Xide Xia, Graham Neubig, Pengchuan Zhang, and Deva Ramanan. Evaluating text-to-visual generation with image-to-text generation. In *Computer Vision – ECCV 2024*, pp. 366–384. Springer Nature Switzerland, 2024. doi: 10.1007/978-3-031-72673-6_20. URL https://doi.org/10.1007/978-3-031-72673-6_20.

Bingshuai Liu, Longyue Wang, Chenyang Lyu, Yong Zhang, Jinsong Su, Shuming Shi, and Zhaopeng Tu. On the cultural gap in text-to-image generation. *arXiv preprint arXiv:2307.02971*, 2023. URL https://arxiv.org/abs/2307.02971.

Sina Malakouti, Boqing Gong, and Adriana Kovashka. Culture in action: Evaluating text-to-image models through social activities. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=opG4m2U0Oo.

Rashid Mushkani, Shravan Nayak, Hugo Berard, Allison Cohen, Shin Koseki, and Hadrien Bertrand. LIVS: A pluralistic alignment dataset for inclusive public spaces. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 45311–45341. PMLR, 2025. URL https://proceedings.mlr.press/v267/mushkani25a.html.

Shravan Nayak, Mehar Bhatia, Xiaofeng Zhang, Verena Rieser, Lisa Anne Hendricks, Sjoerd van Steenkiste, Yash Goyal, Karolina Stanczak, and Aishwarya Agrawal. CulturalFrames: Assessing cultural expectation alignment in text-to-image models and evaluation metrics. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp. 20918–20953. Association for Computational Linguistics, 2025. doi: 10.18653/v1/2025.findings-emnlp.1141. URL https://aclanthology.org/2025.findings-emnlp.1141/.

Aniket Rege, Zinnia Nie, Mahesh Ramesh, Unmesh Raskar, Zhuoran Yu, Aditya Kusupati, Yong Jae Lee, and Ramya Korlakai Vinayak. CuRe: Cultural gaps in the long tail of text-to-image systems. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 15680–15691, October 2025. URL https://openaccess.thecvf.com/content/ICCV2025/html/Rege_CuRe_Cultural_Gaps_in_the_Long_Tail_of_Text-to-Image_Systems_ICCV_2025_paper.html.

Michael Saxon and William Yang Wang. Multilingual conceptual coverage in text-to-image models. *arXiv preprint arXiv:2306.01735*, 2023. URL https://arxiv.org/abs/2306.01735.

Hua Shen. Bidirectional human–AI alignment in education for trustworthy learning environments. *arXiv preprint arXiv:2512.21552*, 2025. doi: 10.48550/arXiv.2512.21552. URL https://arxiv.org/abs/2512.21552.

Hua Shen, Tiffany Knearem, Divy Thakkar, Pat Pataranutaporn, Anoop Sinha, Yike Shi, Jenny T. Liang, Lama Ahmad, Tanu Mitra, Brad A. Myers, and Yang Li. Human–AI interaction alignment: Designing, evaluating, and evolving value-centered AI for reciprocal human–AI futures. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 2026. doi: 10.1145/3772363.3778710. URL https://doi.org/10.1145/3772363.3778710.

Chuancheng Shi, Shangze Li, Shiming Guo, Simiao Xie, Wenhua Wu, Jingtong Dou, Chao Wu, Canran Xiao, Cong Wang, Zifeng Cheng, Fei Shen, and Tat-Seng Chua. Where culture fades: Revealing the cultural gap in text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14702–14712, June 2026. URL https://openaccess.thecvf.com/content/CVPR2026/papers/Shi_Where_Culture_Fades_Revealing_the_Cultural_Gap_in_Text-to-Image_Generation_CVPR_2026_paper.pdf.

Xiujie Song, Mengyue Wu, Kenny Q. Zhu, Chunhao Zhang, and Yanyi Chen. A cognitive evaluation benchmark of image reasoning and description for large vision-language models. In Luis

Chiruzzo, Alan Ritter, and Lu Wang (eds.), *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 6392–6409, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.naacl-long.324. URL https://aclanthology.org/2025.naacl-long.324/.

The Culture Factor. Country comparison tool. https://www.theculturefactor.com/country-comparison-tool, n.d.a. Country-score reference incorporating the October 16, 2023 IDV and LTO update; update provenance documented in the official FAQ. Accessed September 24, 2026.

The Culture Factor. Frequently asked questions. https://www.theculturefactor.com/frequently-asked-questions, n.d.b. Sections on renaming MAS and the October 16, 2023 IDV and LTO score update. Accessed September 24, 2026.

Aleksandra Urman, Elsa Lichtenegger, Salima Jaoua, Azza Bouleimen, Robin Forsberg, Corinna Hertweck, Stefania Ionescu, Nicolò Pagan, Ancsa Hannak, and Joachim Baumann. Prompt revision as a source of cultural bias in text-to-image systems. *arXiv preprint arXiv:2609.11532*, 2026. URL https://arxiv.org/abs/2609.11532.

Mor Ventura, Eyal Ben-David, Anna Korhonen, and Roi Reichart. Navigating cultural chasms: Exploring and unlocking the cultural POV of text-to-image models. *Transactions of the Association for Computational Linguistics*, 13:142–166, 2025. doi: 10.1162/tacl_a_00732. URL https://aclanthology.org/2025.tacl-1.10/.

Yuhang Wang, Yanxu Zhu, Chao Kong, Shuyu Wei, Xiaoyuan Yi, Xing Xie, and Jitao Sang. CDEval: A benchmark for measuring the cultural dimensions of large language models. *arXiv preprint arXiv:2311.16421*, 2023. URL https://arxiv.org/abs/2311.16421.

Yupeng Xie, Zhiyang Zhang, Yifan Wu, Sirong Lu, Jiayi Zhang, Zhaoyang Yu, Jinlin Wang, Sirui Hong, Bang Liu, Chenglin Wu, and Yuyu Luo. VisJudge-Bench: Aesthetics and quality assessment of visualizations. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=lG1HWWdEbN.

Arnav Yayavaram, Siddharth Yayavaram, Simran Khanuja, Michael Saxon, and Graham Neubig. CAIRE: Cultural attribution of images with retrieval. In Vera Demberg, Kentaro Inui, and Lluís Marquez (eds.), *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 8320–8338, Rabat, Morocco, March 2026. Association for Computational Linguistics. doi: 10.18653/v1/2026.eacl-long.389. URL https://aclanthology.org/2026.eacl-long.389/.

Akhila Yerukola, Saadia Gabriel, Nanyun Peng, and Maarten Sap. Mind the gesture: Evaluating AI sensitivity to culturally offensive non-verbal gestures. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 25041–25080, Vienna, Austria, July 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.acl-long.1218. URL https://aclanthology.org/2025.acl-long.1218/.

APPENDIX GUIDE

The appendices follow the benchmark, measurement, validation, and results, followed by prompts, annotation materials, and limitations.

- Appendix A: Benchmark construction and filtering.
- Appendix B: Experimental settings and measures.
- Appendix C: Human validation and annotation diagnostics.
- Appendix D: Extended results.
- Appendix E: Robustness and sensitivity analyses.
- Appendix F: Qualitative examples and failure cases.
- Appendix G: Prompts and annotation materials.
- Appendix H: Limitations and future directions.

A BENCHMARK CONSTRUCTION AND FILTERING

A.1 VALUE DIMENSIONS AND ORIENTATIONS

We use Hofstede's six cultural value dimensions (Hofstede, 2011) (Table 1). For MAS we adopt The Culture Factor's name, *Motivation towards Achievement and Success*, which keeps the dimension's interpretation and country scores (The Culture Factor, n.d.b). Orientation A is the higher-score direction of the country reference; Appendix G.5 gives the annotation criteria.

Table 1: Six value dimensions and their image-level operationalization. The questions summarize the authors' criteria, informed by Hofstede and The Culture Factor.

Code	Value dimension	Orientations (A / B)	Question asked of the depicted situation
PDI	Power distance	Acceptance of unequal authority / equal standing and influence	How are authority and the standing of those affected treated?
IDV	Individualism–collectivism	Personal autonomy / in-group obligation	Whose choices and responsibilities organize conduct?
MAS	Motivation towards achievement and success	Achievement and competition / care and quality of life	What is treated as worth striving for or maintaining?
UAI	Uncertainty avoidance	Need for certainty / tolerance of ambiguity	How is unresolved uncertainty handled?
LTO	Long-term orientation	Perseverance and pragmatic adaptation / quicker returns and normative continuity	How are the timing of returns and established practices treated?
IVR	Indulgence–restraint	Relatively free enjoyment / restraint by norms and duties	How are enjoyment and expression accommodated or constrained?

Country comparisons use The Culture Factor's scores after its October 16, 2023 update to IDV and LTO (The Culture Factor, n.d.a; n.d.b) (Table 2).

Table 2: Country reference scores (The Culture Factor, 2023 update; retrieved September 24, 2026). Higher values correspond to orientation A.

Country	PDI	IDV	MAS	UAI	LTO	IVR
United States	40	60	62	46	50	68
China	80	43	66	30	77	24
India	77	24	56	40	51	26
Thailand	64	19	34	64	67	45

These scores describe societies, whereas S describes evidence in generated images. We therefore read cue shifts against the difference between the two cued countries, even when both lie on the same side of the scale midpoint.

A.2 Comparison with Related Benchmarks

A.3 Scenario Construction and Filtering

We derive scenarios from the 2,953 CDEval questions, which span six dimensions and seven domains. Concept extraction keeps the issue shared by a question and its two options; scenario generation makes that issue depictable while leaving its realization open (Figure 9; prompts in Appendix G.2).

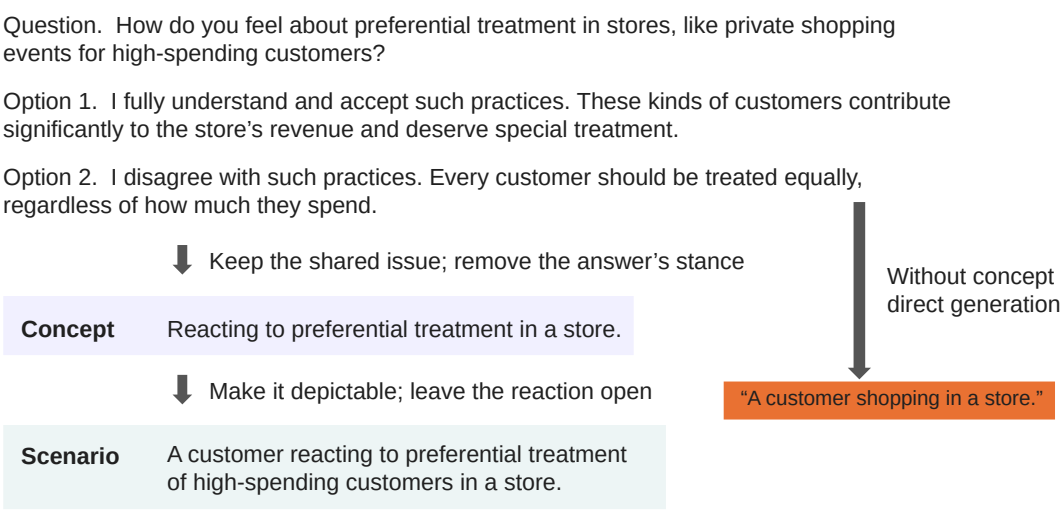

Figure 9: Case study of scenario generation with and without concept extraction using CDEval item `pdi_0052` (Lifestyle) (Wang et al., 2023).

Filtering rules and procedure. We screen all 2,953 scenarios against all six dimensions, not only the source dimension. A scenario is flagged only when quoted wording already fixes an orientation, so that a faithful image could not express the opposite one; for example, "a person indulging in their favorite dessert" prescribes indulgence. Gemini 3.7 Flash and Claude Opus 5 screen each scenario, and we review the union of their flags. Unspecified choices and reactions, roles, objects, or traditional and modern practices shown together do not fix an orientation.

Filtering removes 180 scenarios (Table 4): half are flagged on their source dimension (eight also on another) and half only on other dimensions. We evaluate 160 of the retained scenarios per source dimension, 960 in total.

Table 3: **Comparison of cultural benchmarks and evaluations for text-to-image systems.** Cultural focus denotes the primary target; expectations, practices and values may overlap.

Benchmark	Cultural focus	Prompt design	Multilingual	Geographic cues	Evaluation method
Cultural fidelity and attribution					
CUBE (2024)	Cultural awareness and diversity	Artifacts / broad concepts	✗	Country / no cue	Human ratings; diversity metrics
CuRe (2025)	Long-tail cultural representativeness	Artifacts with varying detail	✗	Region / no cue	Embedding / perceptual metrics; human ratings
CultDiff (2025)	Cultural fidelity	Architecture, clothing, food	✗	Country	Human ratings; learned similarity
CAIRE (2026)	Cultural relevance	Universal cultural concepts	✗	Country	VLM-as-a-judge with retrieval
Cultural representation and conditioning					
ViSAGe (2024)	Visual stereotypes	National identities and attributes	✗	Nationality	Human labels; embedding / caption metrics
Nationalities audit (2025 / rev. 2026)	Attire stereotypes in activities	Everyday activities	✗	Nationality	Human labels; embedding / VLM alignment
CultureBench (2026)	Cultural consistency	Descriptions with / without cultural modifiers	✓	Optional cultural modifiers	VLM cultural attribution
SoS (2026)	Language vs. requested culture	Cultural identity templates	✓	Cultural identity	Image-embedding similarity
WORLDVIEW (2026)	Revision-induced cultural bias	Shared everyday topics	✓	Geographic context / no cue	Embedding metrics; lexical analysis
Cultural expectations, practices and values					
CulturalFrames (2025)	Explicit / implicit expectations	Social practices and scenes	✗	Country context	Human ratings
CULTIVate (2026)	Cultural faithfulness in social activities	Social activities and interactions	✗	Country	Descriptor-based metrics
MC-SIGNS (2025)	Gesture meaning and offensiveness	Gestures and social intentions	✗	Country context	Rejection rates; VLM gesture checks
LIVS (Mushkani et al., 2025)	Community value preferences	Public-space scenes	✗	Local places	Human pairwise preferences
CulText2I (2025)	Cultural axes and associations	Short concept templates	✓	Nationality / no cue	Embedding metrics; VQA
Ours	Cultural value orientations	Social scenarios; value relevance with orientation openness; visual expressibility	✓	Country / no cue	VLM-as-a-judge; evidence proposal and verification

Notes. Multilingual: ✓ indicates reported generation experiments in multiple languages; ✗ indicates no such experiment reported. Retrieval and annotator languages do not count. Geographic cues summarize reported settings, not necessarily a full factorial design. Evaluation lists main approaches; auxiliary validation is omitted. VQA denotes visual question answering. CAIRE refers to its universal-concept generated-image evaluation; MC-SIGNS to its T2I branch.

Table 4: Scenario counts by source dimension before and after filtering.

	PDI	IDV	MAS	UAI	LTO	IVR
Initial	512	472	452	530	485	502
Filtered	25	28	44	27	25	31
Retained	487	444	408	503	460	471

B Experimental Settings and Measures

B.1 Image-generation Systems and Prompt Processing

Table 5: Six image-generation models under eight evaluation configurations. Condition counts refer to benchmark prompts before any system-level prompt processing.

Model	Prompt processing	Conditions	Output resolution
Commercial API systems			
Nano Banana 2	Built-in thinking	14	1024 × 1024
GPT-Image-2	Internal processing	14	Auto
MAI-Image-2.6-Flash	Internal processing	14	1024 × 1024
Open-weight image models			
HiDream-O1-Image	Without Prompt Refiner	6	2048 × 2048
HiDream-O1-Image	With Prompt Refiner	6	2048 × 2048
FLUX.2-dev	Without Prompt Upsampler	6	1024 × 1024
FLUX.2-dev	With Prompt Upsampler	6	1024 × 1024
ideogram-4 (FP8)	Magic Prompt API, enabled	6	1024 × 1024

The main analysis uses the prompt-processed HiDream and FLUX configurations (Table 5); their unrefined counterparts serve as paired controls (Appendix D.4).

Language coverage. In preliminary tests, the open-weight models followed Hindi and Thai prompts less reliably than English and Chinese ones (Figure 10). We therefore evaluate them only under the six shared English and Chinese conditions; the three commercial systems cover all 14 conditions (Table 6).

Scenario: A married couple spending their leisure time.



(a) EN (b) ZH (c) TH (d) HI

Figure 10: **Scenario adherence in a multilingual HiDream pilot.** English and Chinese prompts yield a couple at leisure, while Thai and Hindi prompts do not. HiDream-O1 without refinement, no cue.

Generation settings and prompt processing. Benchmark prompts are submitted unchanged; any rewriting belongs to the system's own pipeline: HiDream's Prompt Refiner, FLUX's Prompt Upsampler, or ideogram's Magic Prompt. We keep original and processed prompts separately, and

Table 6: Configured language and geographic-cue conditions in VisionValueBench. ✓ marks a configured condition and a dash an unconfigured one; * marks the six conditions shared by every configuration, including open-weight generators.

Prompt language	No cue	US cue	China cue	Thailand cue	India cue
English	✓*	✓*	✓*	✓	✓
Chinese	✓*	✓*	✓*	–	–
Thai	✓	✓	–	✓	–
Hindi	✓	✓	–	–	✓

failed HiDream refinements do not fall back to the unrefined prompt. Local generations use seed 20260728 (Table 7).

Table 7: Generation and prompt-processing settings. "Unspecified" marks parameters absent from the recorded API requests.

Component	Model identifier	Settings
Nano Banana 2	`gemini-3.1-flash-image`	Minimal thinking; seed 20260728
GPT-Image-2	`gpt-image-2`	Auto size
MAI	`MAI-Image-2.6-Flash`	$n = 1$; seed unspecified
HiDream	`HiDream-O1-Image-Dev-2604`	With Refiner: flash scheduler, 28 image steps
HiDream Refiner	`HiDream-ai/Prompt-Refine`	$T = 1$, top-$p = .95$, top-$k = 64$; 2,048 output tokens
FLUX	`black-forest-labs/FLUX.2-dev`	50 steps; guidance 4; bfloat16
FLUX Upsampler	`Mistral-Small-3.2-24B-Instruct-2506`	$T = .15$; 512 output tokens
ideogram Magic	`ideogram-4-fp8`	Magic Prompt enabled; 48 steps

B.2 Evaluation Sample and Completion

Every image is evaluated on all six dimensions. We retain one image per available scenario–condition–configuration with complete SAE and IAE judgments from all three verifiers (Table 8); the six main configurations contribute 34,377 images under the shared conditions.

Table 8: Completed image annotations / planned images. Shared conditions are EN/ZH × no/US/China cue; only the three commercial systems have the remaining conditions.

Configuration	Shared conditions	All configured conditions
MAI	5,701 / 5,760	13,302 / 13,440
GPT-Image-2	5,650 / 5,760	13,243 / 13,440
Nano Banana 2	5,753 / 5,760	13,426 / 13,440
FLUX + Upsampler	5,760 / 5,760	5,760 / 5,760
HiDream + Refiner	5,760 / 5,760	5,760 / 5,760
ideogram-4 Magic	5,753 / 5,760	5,753 / 5,760
FLUX control	5,760 / 5,760	5,760 / 5,760
HiDream control	5,760 / 5,760	5,760 / 5,760

Table 9 gives the evaluated images N and the supported-image counts U that serve as score denominators.

B.3 Four-state Outcomes and Directional Scores

Separate support judgments for orientations A and B yield four states per dimension: A only, B only, both, and neither. Neither means that support was not established, not that the image is value-neutral.

Table 9: Supported-image counts U (score denominators) under the six shared conditions. N counts completely annotated images and is shared across dimensions and methods.

Configuration	N	PDI	IDV	MAS	UAI	LTO	IVR
SAE							
MAI	5701	927	3165	2302	674	1650	923
GPT-Image-2	5650	897	2443	1911	631	1242	800
Nano Banana 2	5753	952	2245	1605	514	963	1017
FLUX + Upsampler	5760	678	1708	1005	350	561	765
HiDream + Refiner	5760	741	1812	1257	364	639	768
ideogram-4 Magic	5753	822	1965	1328	410	794	755
FLUX (control)	5760	661	1290	791	234	316	607
HiDream (control)	5760	688	1400	1092	276	433	1044
IAE							
MAI	5701	1827	4896	5384	2252	4558	2013
GPT-Image-2	5650	1743	4052	4916	2026	3605	1960
Nano Banana 2	5753	1644	3463	4526	1328	2573	1938
FLUX + Upsampler	5760	1318	2850	3935	1131	1766	1681
HiDream + Refiner	5760	1311	2827	4126	1094	1870	1690
ideogram-4 Magic	5753	1552	3271	4311	1199	2403	1790
FLUX (control)	5760	1188	2395	3364	905	1336	1429
HiDream (control)	5760	1240	2642	4055	790	1477	2136

Alongside S (Equation 1), we report support prevalence,

$$C = \frac{n_{\text{A only}} + n_{\text{B only}} + n_{\text{both}}}{N}, \tag{2}$$

where N counts images in all four states. S measures direction among supported images and C how often support is established; excluding neither from the denominator keeps sparse evidence from pulling S toward zero. Images supporting both orientations enter S with weight 0. Identical S and C can still arise from orientations co-occurring within images or opposing each other across images.

B.4 Aggregation, Pairing, and Uncertainty

Profiles pool state counts over the six shared conditions before computing S. Language and cue comparisons pair scenarios within each system, fix the other prompt factor, and keep only scenarios available in both arms; each arm keeps its own supported-image denominator. We report $\Delta S = S_{\text{target}} - S_{\text{baseline}}$.

Uncertainty comes from 5,000 stratified scenario-cluster bootstrap replicates. Each replicate resamples 160 source IDs within each source stratum, carrying a scenario's conditions, systems, dimensions and verifier judgments together, and recomputes both arm scores and their difference on each paired cohort. With one image per prompt setting, the intervals reflect variation across scenarios. Pointwise 95% intervals use bootstrap percentiles.[2]

Approximate 95% simultaneous bands use the 95th percentile of the bootstrap maximum standardized absolute deviation. Families are defined per method: the 36 profile cells; the available main configurations $\times$ six dimensions for each language or cue contrast; and the 12 cells of each refinement setting, with unrefined controls kept separate.

C Human Validation and Annotation Diagnostics

C.1 Human Evidence-review Protocol

Five undergraduate and doctoral students in our laboratory, with cultural backgrounds spanning China, India and the United States, reviewed the same 20 images after training and practice. They

[2]The bootstrap seed is 20260925. Intervals are reported only when at least 95% of replicates yield defined estimates.

participated voluntarily without compensation, gave informed consent, and could ask the authors about the instructions (Appendix G.6). Reviewers saw the frozen bank of 388 SAE and 234 IAE candidates from the three production proposers, without proposer identities or automatic decisions, and could add their own evidence.

A reviewer supports an orientation through an accepted candidate or complete added evidence; the human majority needs three of five reviewers. Automatic labels are recomputed from each verifier's accepted candidates in the same bank plus its added evidence, then aggregated two-of-three, so agreement measures evidence review given this bank.

Sample selection. The development cohort is separate from production. After excluding 54 previously exposed lineages, 186 of 240 pooled scenarios were eligible. The sample has five Nano Banana 2 images and three each from HiDream, HiDream with Refiner, FLUX, FLUX with Upsampler and ideogram-4 Magic (no GPT or MAI), drawn from four PDI, four IVR and three scenarios of each other dimension. Two images lacking value content were replaced within the same system, source and condition slot.[3]

C.2 Agreement with Human Review

Table 10 reports agreement and Cohen's κ against each method's own human-majority reference, with intervals from 10,000 image-level bootstrap resamples (seed 20260925). Over the 240 orientation pairs, SAE has 19 positive agreements, 54 human-only positives, no model-only positives and 167 negative agreements; IAE has 52, 45, 2 and 141.

Table 10: Orientation agreement (%) on the 20-image cohort, with 95% intervals on the 0–1 scale. Each dimension has 40 judgments.

	SAE versus SAE-H			IAE versus IAE-H		
Dim.	Agree.	95% CI	κ	Agree.	95% CI	κ
All	77.5	[0.717, 0.829]	0.329	80.4	[0.738, 0.871]	0.562
PDI	82.5	[0.725, 0.925]	0.307	75.0	[0.575, 0.900]	0.322
IDV	55.0	[0.425, 0.675]	0.137	67.5	[0.525, 0.825]	0.387
MAS	75.0	[0.625, 0.875]	0.342	80.0	[0.675, 0.925]	0.614
UAI	75.0	[0.625, 0.875]	0.000	87.5	[0.775, 0.950]	0.600
LTO	87.5	[0.775, 0.950]	0.635	85.0	[0.750, 0.950]	0.652
IVR	90.0	[0.800, 0.975]	0.615	87.5	[0.775, 0.975]	0.684

Positive/negative agreement is 41.3%/86.1% for SAE and 68.9%/85.7% for IAE.[4] Mean pairwise human agreement is 78.58% for SAE and 87.25% for IAE; Fleiss' κ is 0.513 and 0.731. Table 11 compares both automatic methods against each human reference.

Table 11: Agreement against both human-majority references and recovery of human-supported UAI judgments (240 pairs overall, 40 for UAI).

Metric	Automatic method	SAE-H	IAE-H
Overall agreement (%)	SAE	77.50	67.50
	IAE	**82.92**	**80.42**
UAI agreement (%)	SAE	75.0	75.0
	IAE	**87.5**	**87.5**
UAI support recovered	SAE	0/10	0/10
	IAE	**5/10**	**5/10**

[3] Initial selection used seed 20260916; replacements used seed 20260917 and selected the first author-accepted candidate.

[4] Positive and negative agreement are $2a/(2a+b+c)$ and $2d/(2d+b+c)$, where a, d count positive and negative agreements and b, c count disagreements.

C.3 Evidence Acceptance, Omissions, and Disagreements

Human majority accepts 216 of 234 IAE candidates. Automatic majority retains 142 (137 human-accepted) and rejects 92 (79 human-accepted), for 64.10% candidate agreement ($\kappa = 0.124$): verification removes some disputed evidence but also rejects many human-endorsed candidates. Orientation agreement is higher because another retained candidate or an added item can still establish support. SAE candidate agreement is 71.39%, with 43 positive agreements, 111 human-only positives and no model-only positives. Reviewers added 14 evidence records across methods; Appendix F illustrates missed and disputed evidence.

Method-specific support. Of the 240 orientation pairs, IAE alone supports 35, both methods 19, and SAE alone none. The IAE human reference corroborates 33 of the 35 IAE-only pairs and all 19 shared pairs, but also supports 45 of the 186 pairs that neither method supports. IAE's additional support is thus usually corroborated in this cohort, while substantial omissions remain.

UAI support. Both methods agree with all 30 human-unsupported UAI judgments. Of the ten human-supported judgments, SAE recovers none and IAE recovers five (Table 11); all five receive majority support under both human references. These results concern orientation support in the selected cohort, not the direction of individual systems' production UAI scores.

D Extended Results

D.1 Complete Value Profiles

Table 12 pools all available images in the six shared conditions for each main configuration. The two methods agree on the sign of every dimension except UAI.

Table 12: Pooled orientation scores under the six shared conditions. Bold: within-method simultaneous 95% interval excludes zero.

Method	System	PDI	IDV	MAS	UAI	LTO	IVR
SAE	MAI	**−0.886**	**0.176**	**−0.420**	**−0.239**	**0.733**	**0.777**
SAE	GPT Image 2	**−0.805**	**0.273**	**−0.250**	−0.124	**0.539**	**0.787**
SAE	Nano Banana 2	**−0.887**	**0.297**	**−0.383**	**−0.276**	**0.487**	**0.885**
SAE	FLUX + Upsampler	**−0.704**	**0.458**	**−0.304**	−0.271	**0.398**	**0.890**
SAE	HiDream + Refiner	**−0.741**	**0.374**	**−0.284**	**−0.275**	**0.376**	**0.919**
SAE	Ideogram Magic	**−0.845**	**0.348**	**−0.345**	−0.237	**0.374**	**0.868**
IAE	MAI	**−0.854**	**0.089**	**−0.267**	−0.042	**0.688**	**0.581**
IAE	GPT Image 2	**−0.740**	**0.196**	**−0.164**	**0.165**	**0.546**	**0.596**
IAE	Nano Banana 2	**−0.785**	**0.256**	**−0.219**	**0.187**	**0.594**	**0.807**
IAE	FLUX + Upsampler	**−0.631**	**0.341**	**−0.206**	0.142	**0.531**	**0.838**
IAE	HiDream + Refiner	**−0.709**	**0.323**	**−0.205**	**0.182**	**0.575**	**0.848**
IAE	Ideogram Magic	**−0.740**	**0.293**	**−0.230**	**0.285**	**0.479**	**0.789**

Why MAI's UAI score is near zero. Under IAE, MAI-2.6-Flash's UAI score is close to zero because its images support the two UAI orientations about equally often. Across the six shared conditions, 935 images support only high UAI, 1,029 support only low UAI, 288 support both, and 3,449 support neither, so $S = (935 - 1{,}029)/(935 + 1{,}029 + 288) \approx -0.042$. Images supporting both orientations enter the denominator but not the numerator, which pulls the score further toward zero.



Figure 11: **Added text supports both UAI orientations.** MAI-2.6-Flash, EN / no cue.

Many of these mixed images come from text that MAI adds to the scene. In Figure 11, the prompt describes a person reacting to an artwork while others criticize it. MAI adds speech bubbles and slogans. A critic says "I don't get it. It looks unfinished," rejecting the painting's ambiguity, while the central woman holds a booklet reading "It

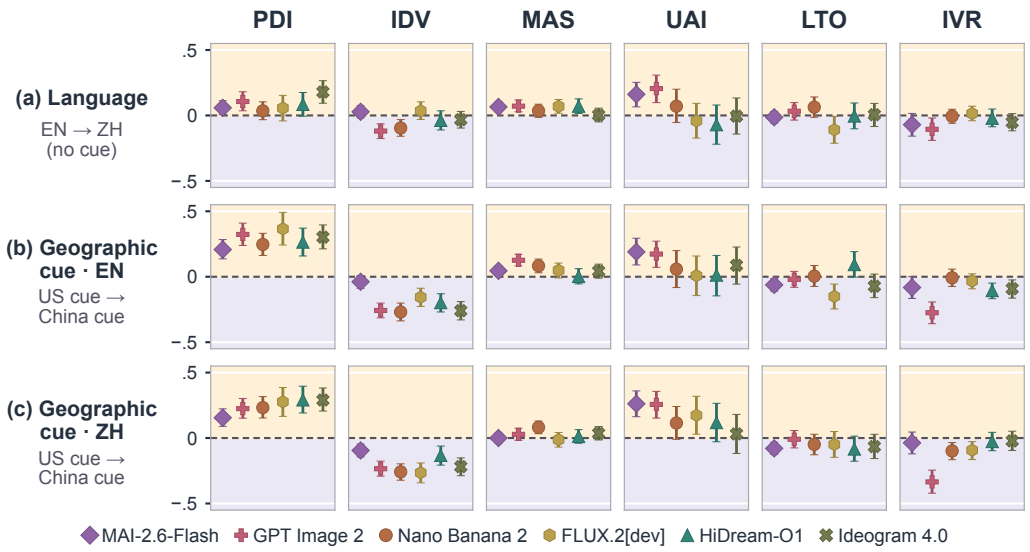

Figure 12: **Under IAE, geographic cues shift scores more than prompt language, especially PDI and IDV.** Paired changes for the comparisons in Figure 7, with pointwise 95% bootstrap intervals.

speaks to me," accepting it. All three IAE verifiers therefore support both high and low UAI for this image.

D.2 Prompt-Language and Geographic-Cue Effects

Table 13 summarizes all English/Chinese comparisons in the six shared conditions (procedure in Appendix B.4). A kept sign and an interval excluding zero answer different questions: whether the direction persists, and whether the score changes.

Table 13: Paired shifts in the six shared conditions: 36 system–dimension comparisons per method, 923–960 paired scenarios per system. Med./Max.: $|\Delta S|$; Keep: unchanged signs; Sim.: simultaneous 95% intervals excluding zero.

	SAE				IAE			
Contrast; fixed condition	Med.	Max.	Keep	Sim.	Med.	Max.	Keep	Sim.
EN → ZH; no cue	0.052	0.242	36	2	0.057	0.205	36	5
EN → ZH; US cue	0.046	0.248	36	0	0.031	0.187	36	1
EN → ZH; China cue	0.049	0.270	36	0	0.035	0.148	36	0
none → US; EN	0.045	0.193	36	0	0.027	0.140	36	1
none → China; EN	0.078	0.372	36	11	0.117	0.337	35	18
US → China; EN	0.074	0.409	36	10	0.092	0.366	35	16
none → US; ZH	0.054	0.252	36	0	0.043	0.190	36	4
none → China; ZH	0.075	0.327	36	3	0.090	0.253	35	9
US → China; ZH	0.111	0.454	36	7	0.097	0.335	35	17

For English US-to-China cues under IAE, PDI rises and IDV falls in every system; all six PDI changes and five IDV changes (all but MAI's) exclude zero under both pointwise and simultaneous intervals. SAE gives the same directions. LTO changes are mixed under IAE and negative in all six systems under SAE, opposite to the country-reference difference (Table 2).

D.3 Effect Magnitudes and No-cue Profiles

Figures 13 and 14 compare no-cue prompts with US- and China-cue prompts, by single-dimension changes and by whole profiles.

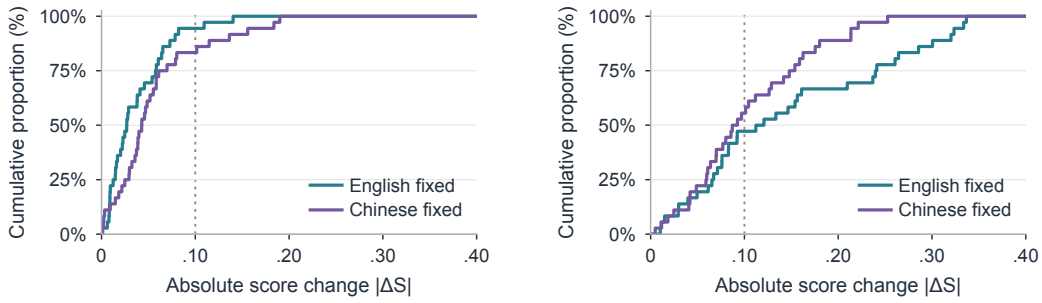

(a) No cue → US cue.

(b) No cue → China cue.

Figure 13: **Adding a China cue changes IAE scores more than adding a US cue.** Each curve covers 36 system–dimension comparisons; the dotted line marks 0.10 (counts in Table 21).

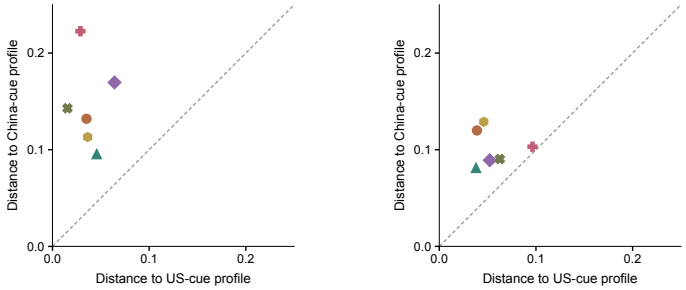

(a) English fixed.

(b) Chinese fixed.

Figure 14: **No-cue profiles are closer to US-cue than to China-cue profiles.** Each point is one system's mean absolute IAE score difference across the six dimensions, on scenarios present in all six conditions (917–960 per system).

D.4 Effects of Prompt Refinement

Table 14: Score changes from prompt refinement, ΔS [pointwise 95% CI], on 960 paired scenarios across six conditions (5,760 images per arm).

Method	Dimension	HiDream Refiner	FLUX Upsampler
SAE	PDI	+0.128 [+0.057, +0.203]	−0.136 [-0.232, -0.046]
SAE	IDV	+0.248 [+0.166, +0.330]	+0.166 [+0.097, +0.234]
SAE	MAS	+0.097 [+0.008, +0.186]	−0.005 [-0.099, +0.088]
SAE	UAI	+0.602 [+0.418, +0.786]	−0.246 [-0.447, -0.041]
SAE	LTO	+0.182 [+0.021, +0.339]	+0.230 [+0.075, +0.385]
SAE	IVR	−0.050 [-0.086, -0.017]	+0.057 [-0.012, +0.129]
IAE	PDI	+0.121 [+0.063, +0.181]	−0.080 [-0.152, -0.012]
IAE	IDV	+0.131 [+0.087, +0.176]	+0.068 [+0.030, +0.106]
IAE	MAS	+0.122 [+0.089, +0.156]	+0.037 [+0.003, +0.071]
IAE	UAI	+0.539 [+0.422, +0.649]	−0.269 [-0.353, -0.189]
IAE	LTO	+0.007 [-0.057, +0.070]	+0.063 [+0.009, +0.120]
IAE	IVR	−0.083 [-0.121, -0.048]	+0.024 [-0.016, +0.068]

The two refiners change scores differently. Under IAE, HiDream's UAI reverses from −0.357 to +0.182, while FLUX's falls from +0.411 to +0.142. Support prevalence rises for every family, method and dimension except HiDream's IVR (18.1% to 13.3% under SAE, 37.1% to 29.3% under IAE). The comparisons capture the whole prompt-processing configuration, not particular rewritten words.

Case study: from a family celebration to a personal pursuit. Refinement can resolve a choice that the original prompt leaves open (Figure 15). The prompt describes "a person following through on the choice they made between attending a family celebration and pursuing a personal plan that brings enjoyment." HiDream's refiner rewrites it as "a person with a joyful expression [who] stands on a mountain summit with arms outstretched," dropping the family alternative. The unrefined image shows a family celebration and the refined one a solo outdoor activity; all three verifiers under both SAE and IAE support only collectivism for the former and only individualism for the latter. We hypothesize that such changes arise jointly from prompt refinement and the generator's interpretation of the rewritten prompt; future work will separate the two.



(a) Without refiner: collectivism only.



(b) With refiner: individualism only.

Figure 15: **HiDream-O1 with and without prompt refinement** (English, no cue; CDEval `ivr_0356`; matched prompt and seed).

D.5 Additional Language and Geographic-Cue Shifts

Figure 16 shows the 27 IAE contrasts beyond the shared-condition comparisons above (EN, ZH, HI, TH: English, Chinese, Hindi, Thai). EN-to-ZH contrasts under US and China cues cover all six systems; contrasts involving Hindi, Thai, India or Thailand cover the three commercial systems.

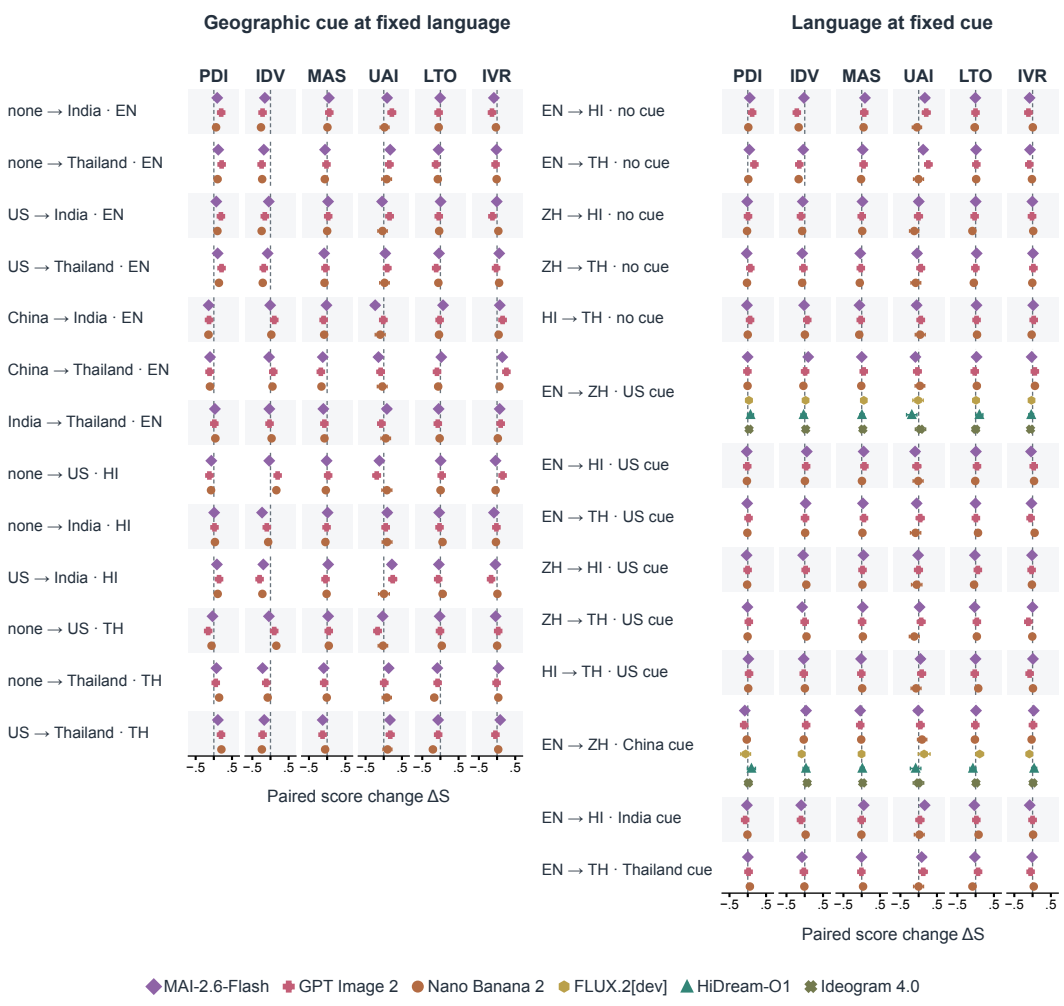

Figure 16: **Additional IAE contrasts.** Left: geographic cue at fixed language; right: language at fixed cue. Rows give baseline → target and the fixed factor. Bars are pointwise 95% bootstrap intervals (923–960 matched scenarios).

E Robustness and Sensitivity Analyses

All analyses use the six-system, shared-condition cohort and vary one choice at a time: the verifier rule, score aggregation, source strata, the effect-magnitude threshold, or the scenario sample.

E.1 Verifier Selection and Aggregation

We recompute each method under any-verifier (1/3), majority (2/3), unanimity (3/3) and single-verifier rules, keeping images and individual judgments fixed (Table 15).

Under SAE, only Gemini-only verification changes a sign (GPT-Image-2's UAI). Under IAE, unanimity and Gemini-only verification turn IDV negative for MAI, Nano Banana 2 and ideogram-4, and any-verifier and Luna-only verification change MAI's UAI sign. Unanimity also cuts IAE support prevalence from 62.1% to 25.4% for IDV and from 79.1% to 42.3% for MAS. Shared directions therefore depend partly on the verifier rule.

Paired language and cue effects. Table 16 repeats two paired comparisons on the same matched scenarios under each rule.

Table 15: Profile sensitivity to the verifier rule over 36 cells per method. Max.: largest $|\Delta S|$ from the 2/3 rule; Rev.: sign changes.

	SAE		IAE	
Rule	Max.	Rev.	Max.	Rev.
2/3	0.000	0	0.000	0
1/3	0.153	0	0.118	1
3/3	0.224	0	0.407	3
Luna	0.210	0	0.199	1
Gemini	0.224	1	0.409	3
Muse	0.114	0	0.136	0

Table 16: Matched-contrast sensitivity for IAE EN→ZH with no cue and SAE US→China cue in English. Keep: unchanged signs out of 36; the median refers to the language contrast; PDI up and IDV down: systems out of six.

Rule	IAE language keep	Median $\|\Delta S\|$	SAE cue keep	PDI up	IDV down
2/3	36	0.057	36	6	6
1/3	35	0.045	36	6	5
3/3	34	0.081	34	6	6
Luna	35	0.052	35	6	5
Gemini	34	0.096	30	6	6
Muse	36	0.046	36	6	6

E.2 Verifier Tendencies and Human Agreement

Within each method the three verifiers see the same candidate bank but differ in acceptance, support and human agreement (Table 17).

Table 17: Verifier diagnostics (%, except κ). Support and acceptance use the production cohort (412,524 image–orientation slots per method; 662,844 SAE and 434,932 IAE candidate appearances per verifier); agreement uses the 240 human-reviewed pairs.

	SAE				**IAE**			
Verifier	Support	Accept	Agree.	κ	Support	Accept	Agree.	κ
Luna	13.0	13.7	78.8	0.375	29.4	69.3	80.8	0.575
Gemini	6.7	9.0	74.6	0.223	13.2	36.7	71.2	0.328
Muse	12.4	14.5	80.0	0.420	30.7	72.1	85.0	0.677

Gemini is the strictest verifier in both methods and agrees least with human review; in IAE it is the sole dissenter in 53,726 of 93,131 split decisions (57.7%). Pairwise verifier agreement is 91.5–93.0% for SAE ($\kappa = 0.526$–0.684) and 81.8–90.4% for IAE ($\kappa = 0.491$–0.773). All three are Western-developed VLMs, so this heterogeneity and the selected human review leave shared cultural bias unresolved.

E.3 Candidate Breadth and Method-specific Support

Table 18 compares both methods on the same 34,377 images. SAE candidates are nearly balanced between orientations, whereas IAE candidate shares vary by dimension. IAE adds the most support for MAS: 23,512 orientation slots are IAE-only, against 274 SAE-only. With an empty candidate bank, SAE supports none of 137,293 orientation slots and IAE six of 217,416, through added evidence.

Table 18: Proposal and support diagnostics on the shared six-system cohort. Candidate shares count orientation-A appearances; support counts use two orientation slots per image (68,754 per dimension).

	PDI	IDV	MAS	UAI	LTO	IVR
SAE A candidates, %	49.1	54.5	48.4	50.1	50.2	52.2
IAE A candidates, %	31.1	58.0	44.8	65.8	72.2	75.1
Supported, SAE only	746	1,998	274	570	728	146
Supported, IAE only	5,413	12,461	23,512	7,302	13,697	6,644
Supported, both	4,276	11,546	9,406	2,381	5,258	4,915
Supported, neither	58,319	42,749	35,562	58,501	49,071	57,049

E.4 Scenario Wording and Source Strata

To test sensitivity to the originating question, we compare scores within each dimension's own source stratum and within the other five (Table 19). Only one SAE UAI cell and one IAE IDV cell change sign.

Table 19: Score ranges across the six systems within each dimension's own source stratum and within the other five strata.

	SAE		IAE	
Dim.	Own source	Other sources	Own source	Other sources
PDI	$[-0.904, -0.711]$	$[-0.930, -0.697]$	$[-0.793, -0.602]$	$[-0.893, -0.653]$
IDV	$[0.066, 0.483]$	$[0.207, 0.448]$	$[-0.021, 0.337]$	$[0.113, 0.343]$
MAS	$[-0.467, -0.194]$	$[-0.406, -0.251]$	$[-0.210, -0.064]$	$[-0.279, -0.175]$
UAI	$[-0.275, 0.062]$	$[-0.437, -0.238]$	$[-0.065, 0.284]$	$[-0.031, 0.286]$
LTO	$[0.457, 0.780]$	$[0.218, 0.711]$	$[0.601, 0.718]$	$[0.418, 0.681]$
IVR	$[0.692, 0.880]$	$[0.832, 0.946]$	$[0.532, 0.800]$	$[0.601, 0.873]$

E.5 Aggregation and Available-condition Coverage

Equal weights over supported scenarios, cohorts with all six conditions present, and a denominator excluding both-supported images all preserve the 36 primary directions (Table 20).

Table 20: Largest $|\Delta S|$ from the primary profile across 36 cells per method under alternative aggregations.

Alternative	SAE	IAE
Equal supported scenarios	0.133	0.173
All six conditions present	0.011	0.010
Both excluded	0.019	0.164

E.6 Sensitivity to the Effect-magnitude Threshold

Table 21 varies the magnitude cutoff τ around the main-text value of 0.10 with the paired estimates fixed; the cutoff describes effect size, while uncertainty comes from the paired intervals.

Under every cutoff most language changes fall below τ, fewer US-to-China cue changes do, and adding a US cue stays mostly small. Whether most cue changes count as small can depend on τ: for SAE US-to-China changes with Chinese prompts, the count below threshold rises from 15/36 at 0.075 to 22/36 at 0.125.

Table 21: Number (percentage) of the 36 system–dimension comparisons with $|\Delta S| < \tau$. Language contrasts fix the cue; cue contrasts fix the language.

		SAE			IAE		
Contrast	Fixed	.075	.100	.125	.075	.100	.125
EN → ZH	No cue	25 (69%)	27 (75%)	31 (86%)	27 (75%)	29 (81%)	33 (92%)
EN → ZH	US cue	26 (72%)	30 (83%)	31 (86%)	32 (89%)	34 (94%)	35 (97%)
EN → ZH	China cue	25 (69%)	27 (75%)	30 (83%)	27 (75%)	33 (92%)	35 (97%)
No → US cue	EN	25 (69%)	30 (83%)	32 (89%)	32 (89%)	34 (94%)	35 (97%)
No → US cue	ZH	22 (61%)	29 (81%)	29 (81%)	28 (78%)	30 (83%)	32 (89%)
No → China cue	EN	18 (50%)	20 (56%)	22 (61%)	11 (31%)	17 (47%)	19 (53%)
No → China cue	ZH	18 (50%)	23 (64%)	26 (72%)	14 (39%)	20 (56%)	23 (64%)
US → China cue	EN	18 (50%)	21 (58%)	21 (58%)	14 (39%)	19 (53%)	20 (56%)
US → China cue	ZH	15 (42%)	17 (47%)	22 (61%)	13 (36%)	19 (53%)	21 (58%)

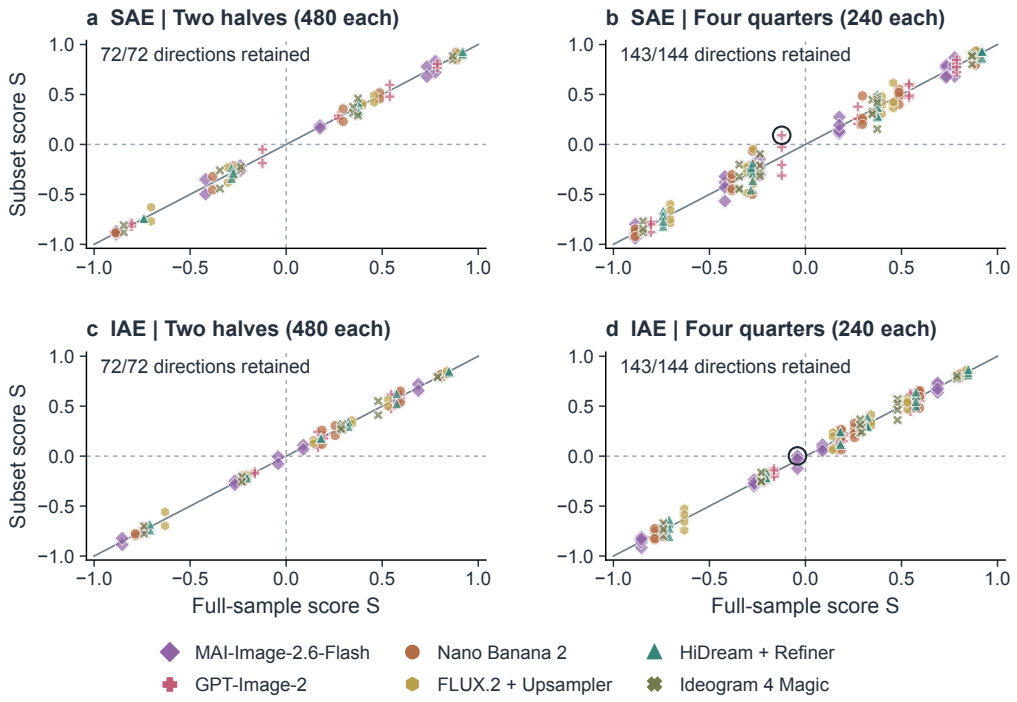

Figure 17: **Value profiles across scenario partitions.** Each point compares a system–dimension score in a half or quarter with the full-sample score; open rings mark the two sign changes. Colors and markers follow Figure 6.

E.7 Stability across Scenario Partitions

We split the 960 scenarios into four quarters of 240 (40 per source dimension) and two halves of 480 (H1=Q1+Q2; H2=Q3+Q4), using a fixed random partition (seed 20260926) chosen without inspecting outcomes that keeps repeated scenario texts together. Profiles and paired effects are recomputed within each group.

Both halves keep all 72 profile directions and the quarters keep 286 of 288 (Figure 17), with largest deviations of 0.097 and 0.226. Both sign changes are in UAI: GPT-Image-2 under SAE (-0.124 to 0.092 in Q1) and MAI under IAE (-0.042 to 0.002 in Q2).



(a) Indulgence only. GPT-Image-2; ZH / China.



(b) Restraint only. GPT-Image-2; ZH / none.



(c) Both IVR orientations. MAI; EN / none.



(d) IDV missed. ideogram-4 Magic; EN / China.



(e) Informed participation. Nano Banana 2; ZH / US.



(f) Co-seating insufficient. FLUX + Upsampler; ZH / US.



(g) Disputed participation. ideogram-4 Magic; EN / none.



(h) Contextual chart. HiDream control; ZH / none.

Figure 18: **Qualitative cases.** (a–c) One windfall scenario realized three ways; (d) a cultural-festival participant; (e–g) low-PDI evidence in a consultation, an elder discussion and a family-tradition scene; (h) a team meeting. Labels give the prompt language and geographic cue.

Prompt effects vary more. Under IAE, 27–32 of 36 language changes per cue stay below 0.10 in the halves and 24–30 in the quarters; US-to-China cues raise PDI in all six systems in every group and lower IDV in all six in each half and five or six in each quarter. SAE is less stable: only 17 of 36 language changes fall below 0.10 in Q3 with no cue and in Q4 with a China cue. Within the benchmark, the broad direction pattern is stable, while finer prompt effects depend more on scenario composition.

F Qualitative Examples and Failure Cases

F.1 Four-state Decisions and Missed Evidence

Panels (a–c) of Figure 18 realize one scenario: "a person following through on the choice they made between using a financial windfall for personal pursuits and investing it in family needs." In the GPT-Image-2 pair, checked personal pursuits, travel and creative activities support indulgence, while pointing toward household obligations and away from musical pursuits supports restraint; all three IAE verifiers agree. The MAI image supports both IVR orientations through different parts of the allocation: family-trip entries accommodate shared enjoyment, while postponed personal luxury supports restraint.

In (d), a festival participant presents their contribution. The human majority supports collectivism (individualism gets two of five votes), but automatic IAE supports neither IDV orientation; humans and IAE both reject both UAI orientations. An automatic "neither" can thus reflect either a shared negative judgment or missed support.

F.2 Interpretation Boundaries and Contextual Evidence

In (e), shared alternatives, risk information and the patient's pen over the decision materials support informed participation under both human and automatic majority. In (f), both majorities reject equal-height seating alone as low-PDI evidence. In (g), the human majority rejects automatic low-PDI support for a child offering a photograph to a note-taking adult: visible participation need not establish influence over a decision.

Background content counts through its role in the situation. In (h), the performance chart supports achievement under both majorities, but automatic review misses the human-endorsed high-UAI reading of the structured information. The family photographs in (g) support enduring in-group continuity under both majorities.

G PROMPTS AND ANNOTATION MATERIALS

This appendix summarizes the prompts and decision criteria in manuscript terminology.

G.1 PIPELINE INDEX

Table 22 indexes construction and annotation calls. Filtering follows in Appendix G.3; annotation rules and human materials in Appendices G.5–G.6.

Table 22: Prompts used to build and annotate VisionValueBench, in pipeline order. Inputs and outputs are summarized; numbers link to the annotation prompts in Appendix G.4.

#	Prompt	Key inputs	Main outputs	Calls
Scenario construction (Section 3)				
–	Concept extraction	CDEval item + metadata	Concept + rationale	1/item
–	Scenario generation	Same as concept extraction + concept	Scenario + rationale	1/concept
–	Translation	English scenario	Chinese, Thai, Hindi scenarios	1/scenario
Annotation (Section 4)				
1	SAE proposal	Scenario	Potential orientation evidence	3/scenario
2	IAE proposal	Image + scenario	Candidate orientation evidence	3/image
3	SAE verification	Image + scenario + pooled candidates from 1	Candidate reviews + 12 orientation decisions	3/image
4	IAE verification	Image + scenario + pooled candidates from 2	Candidate reviews + added evidence + 12 orientation decisions	3/image

Inputs. A CDEval item contains a question and two options; metadata includes its domain, dimension, and orientation descriptions. Highlighted fields such as `{{SCENARIO}}` are filled at run time.

Models and calls. Concept extraction and scenario generation use DeepSeek-V4-Flash; translation uses Gemini 3.7 Flash. Each annotation prompt runs independently on GPT-5.6 Luna, Gemini 3.7 Flash, and Muse Spark 1.3 (three calls in total). Each call uses one user message; in Prompts 2–4, the image follows the text.

G.2 CONSTRUCTION, TRANSLATION, AND RENDERING PROMPTS

Both construction calls use DeepSeek-V4-Flash at temperature 0; their rationales are never sent to image generators.

Construction prompts. *Concept extraction* reads a CDEval question, its two options, domain, source dimension and orientation descriptions, and returns one concise English phrase naming their shared situation, decision or issue. It preserves the source context without preferring either option or adding country, language or cultural-group detail. *Scenario generation* turns the concept into a natural, culture-neutral, image-depictable scenario that requires a reaction or decision while leaving its direction open, so that different orientations yield visibly different realizations of one still image rather than differences in text or subtle details; few-shot outputs include "a person with the car they chose to buy." *Translation* renders each finalized English scenario jointly into Chinese, Thai and Hindi, preserving detail, event status, alternatives and deliberate omissions without cultural adaptation (Gemini 3.7 Flash, medium thinking, 4,096-token output limit).

Prompt rendering. Rendering uses no model: the language's base prompt wraps the scenario `{{SCENARIO}}`, and a geographic cue appends to the rendered prompt `{{PROMPT}}` one sentence naming the country `{{CULTURE}}` (Table 23).

Table 23: Prompt rendering templates. Thai drops terminal punctuation before the cue.

Language	Base prompt	Geographic-cue template	Country labels
English	`Generate an image: {{SCENARIO}}`	`{{PROMPT}}. Set the scene in {{CULTURE}}.`	the United States China Thailand India
Chinese	`生成一张图片：{{SCENARIO}}`	`{{PROMPT}}。将场景设在{{CULTURE}}。`	中国 美国
Thai	`สร้างรูปภาพ: {{SCENARIO}}`	`{{PROMPT}} โดยให้ฉากอยู่ใน{{CULTURE}}`	ประเทศไทย สหรัฐอเมริกา
Hindi	`एक चित्र बनाएँ: {{SCENARIO}}`	`{{PROMPT}}। दृश्य को {{CULTURE}} में रखें।`	भारत संयुक्त राज्य अमेरिका

G.3 Scenario Filtering Prompt

The screening prompt is one user message with the task instructions, the attribution rules and decision cards for all six dimensions, an output contract and the scenario. Its task instructions read:

> The sentence at the end of this message is a prompt for a text-to-image model. The model will draw the scene it describes, and how the scene is drawn can express a cultural value orientation. For each of the six value dimensions in the rubric below, decide whether the sentence already commits the image to one orientation, so that any image faithful to the sentence has to express it.

For each dimension, the model decides whether the sentence involves it and whether its stated conduct commits a faithful image to an orientation; unresolved choices or responses stay open, and a settled orientation must be justified by a quote from the sentence. A validator checks dimension order, orientation keys and that each settling quote occurs in the sentence. Appendix A.3 describes the human review of flags.

G.4 Evaluation Prompts

All four prompts enforce JSON schemas and run independently on GPT-5.6 Luna, Gemini 3.7 Flash and Muse Spark 1.3 with medium reasoning effort. Output limits are 8,192 tokens for Luna and Gemini (32,768 for Gemini completion runs) and 32,768 for Muse.

Shared proposal rules. Rank all six dimensions by relevance (ranks 1–6, ties allowed). A dimension is related only if its decision card's defining issue occurs in the focal situation; relevance does not imply orientation support. Propose zero or more distinct evidence candidates per orientation, with empty lists for unrelated dimensions, fair consideration of both orientations and no target candidate count. Each candidate includes an identifier, an observable basis and its interpretation under the card. Within each method, pool the three models' candidates by orientation, merging only identical candidates.

Prompt 1: SAE proposal — summary: scenario only

Identify the focal situation, what the scenario establishes and what it leaves open. Rank dimension relevance without adding backstory, then describe possible observable depictions and how each could support an orientation. No image, generator identity, country or preferred result is assumed. Reject object inventories, paraphrases and interpretations supported only by an element's presence or text naming an orientation.

Return `focal_situation`, `ranking` and `dimensions`; candidate fields are `id`, `potential_evidence` and `interpretation`. These possibilities are reused for every image generated from the scenario.

Prompt 2: IAE proposal — summary: image + scenario

Rank relevance to the depicted focal situation, then propose candidates grounded in the image. The scenario supplies context; the image determines what is present. Explain each visible basis, its contextual role and its orientation connection. Revise or drop candidates that merely label an isolated element or change the matter being judged; rejecting one orientation does not establish its opposite.

Return `ranking` and `dimensions`; candidate fields are `id`, `visual_basis` and `interpretation`. These are proposals, not final orientation decisions.

Shared verification rules. Verifiers receive the image, scenario and pooled candidates, with identifiers `C1`, `C2`, etc., but no proposer identities, relevance rankings or focal-situation summaries. First describe the image independently of the candidates. Then review every candidate once: identify its core scene; check whether that scene is `grounded`, `not_grounded` or `unclear`; when grounded, test whether its interpretation is `warranted`, `not_warranted` or `unclear` under the card. Otherwise interpretation is `not_assessed`. Accept exactly when both grounding and interpretation pass. Incidental visual differences do not invalidate a matching core scene. Inspect the image throughout, judge candidates as written, assess repeated content consistently and impose no target acceptance rate.

Prompt 3: SAE verification — summary: image + scenario + pool

Check the scenario-derived possibilities against what the image actually depicts. For each of the twelve orientations, return `supported` exactly when at least one candidate is accepted; otherwise return `not_supported`. No additional evidence is introduced. Explain the accepted candidate's relation to the card or the decisive evidential gap. Candidate frequency and agreement do not justify support.

Prompt 4: IAE verification — summary: image + scenario + pool

Apply the same grounding and interpretation checks to image-derived candidates. After reviewing an orientation's candidates, add any missed evidence under the same standard, with `visual_basis` and `interpretation`; otherwise leave the list empty. A claim supported for a different reason than the candidate gives belongs here. Return `supported` exactly when a candidate is accepted or valid additional evidence is recorded.

Shared output contract. Both verification outputs give, for each of the twelve dimension–orientation pairs, the candidate reviews, a decision and an explanation; IAE adds its additional evidence. The schema enforces complete coverage and decisions consistent with the rules above. "Not supported" means insufficient support, not evidence for the opposite orientation.

G.5 Decision Cards for Automatic Annotation

The four annotation prompts share identical dimension cards, condensed in Table 24. All cards share five attribution rules. *Scope:* the scenario supplies context such as alternatives, the object of a response and an arrangement's purpose, while the image decides any open outcome and prevails if the two conflict. *Evidence:* support connects the depiction, its role in the focal situation and the card's criterion; readable text counts in context, but nationality, ethnicity, gender and cultural markers alone never do. Each orientation needs its own positive support, without proof of private motives or a depicted opposite. *Choices:* an indicated choice or a clearly realized alternative can suffice; merely displaying alternatives cannot. *Responses and arrangements:* a depicted practice can show what someone responds to without showing their stance, unless the arrangement itself is the focal matter. *Applicability:* each card first identifies its defining issue in the focal situation; topical association alone does not establish relevance.

Table 24: Condensed decision cards for automatic annotation. The last column lists cues that do not settle the dimension on their own.

Dim.	Defining issue	Orientation criteria	Insufficient alone
PDI	Authority or control over decisions, resources or participation, and how unequal power is treated	*High:* higher standing entitles control or deference, and affected people accept limits on their voice. *Low:* their standing to contribute, question or seek accountability is recognized, even if a leader keeps final authority.	Following instructions (for High); disagreement, informality or friendliness (for Low)
IDV	A personal choice, goal, responsibility or belonging relative to a lasting in-group	*Individualism:* the matter is the person's own to direct, including privacy and self-reliance. *Collectivism:* membership, harmony, loyalty or mutual commitments shape the choice, responsibility or identity.	Separate or shared activity, which may be task allocation or coordination; help and group ties
MAS	What the depiction treats as worth pursuing or maintaining	*Achievement:* accomplishment, advancement, competitive standing or material success is valued. *Care:* needs, well-being, relationships or quality of life matter.	Working, helping or holding a care role; effort, wealth, mood or gender composition
UAI	What is uncertain, ambiguous or unfamiliar for the participants	*High:* unease with uncertainty, or definite answers and fixed procedures treated as needed assurance. *Low:* ambiguity, provisional answers or flexibility accepted while something remains unsettled.	Error prevention, information gathering, precautions against known hazards
LTO	Present investment and expected returns, or established practices and expectations	*Long-term:* commitments serve returns worth waiting for, or a practice is pragmatically adapted. *Short-term:* quicker returns guide effort, or traditions and established expectations guide conduct.	Future benefits, scheduling or fast completion; old and new elements shown together
IVR	The place given to enjoying life, leisure or fun	*Indulgence:* enjoyment is freely pursued, welcomed or accommodated. *Restraint:* enjoyment is curbed or subordinated to norms or duties.	Entertainment, consumption, smiles or luxury; practical limits such as budgets or safety



Prompt: a teacher and students communicating in a classroom.

1. Identify the power relationship. The teacher has authority; the students are less powerful. The matter is a student's opportunity to speak and be heard.

2. Examine its treatment. A student speaks and gestures while the teacher leans toward her and listens with a notepad. The response treats her contribution as worth hearing and considering.

Interpretation: Low PDI. The student's voice is welcomed within the teacher–student relationship even though the teacher retains authority.

Figure 19: The human-review PDI worked example, condensed.

G.6 Decision Cards for Human Review

Reviewers use plain-language versions of the same six cards on the review page, each with a worked two-step example; Figure 19 shows the PDI example.

H Limitations and Future Directions

From cultural constructs to visual evidence. Our evaluation operationalizes cultural values through Hofstede's six dimensions and their associated visual criteria. Still images provide limited access to motives, long-term commitments, and broader social relationships. The same depicted behavior may also support different interpretations across cultural contexts. Our findings therefore concern expressed value orientations under the adopted rubric. Cultural researchers and communities could validate these mappings and compare alternative value frameworks.

Dependence on the scenario distribution. The observed value profiles depend on the benchmark's scenario distribution. Leaving value orientations unspecified does not ensure that opposing orientations are equally easy to depict or recognize. Translating a shared scenario bank supports controlled comparisons across prompt languages but provides limited coverage of locally specific situations. Community-authored scenarios could test whether the observed homogeneity persists across social contexts.

Evaluator dependence and validation scope. Evidence proposal and verification rely on VLMs, which may share omissions or interpretive biases. Human validation involved five laboratory reviewers and 20 development images, using model-proposed candidates with opportunities to add evidence. This design measures agreement primarily within a shared candidate bank and provides limited evidence about the completeness of evidence proposal. Broader validation across reviewers, scenarios, and T2I systems could assess generalizability.

System-level evaluation and component attribution. We evaluate each T2I system as a whole, including its associated prompt-processing components where applicable. Our analysis does not disentangle the contributions of these components and the underlying image-generation model to the expressed value orientations. Future work could characterize their respective contributions and interactions to better understand how value orientations emerge throughout the generation pipeline.

Measurement resolution and downstream use. Binary support judgments record whether evidence supports an orientation but do not measure its visual prominence. Aggregate scores can conceal differences between orientations co-occurring within images and opposing orientations appearing across images. Retaining one image per available scenario, condition, and configuration leaves repeated-generation variability unmeasured, and missing outputs limit the matched comparison cohorts. Whether users perceive these orientations or are influenced by them requires user studies.