

Evaluating AI-Generated Images of Cultural Artifacts with Community-Informed Rubrics

NARI JOHNSON, Microsoft Research, UK and Carnegie Mellon University, USA

DEEPTHI SUDHARSAN*[†], RiskSpan, India

HAMNA*, Microsoft Research, India

SAMANTHA DALAL, Microsoft Research, UK and Princeton University, USA

THEO HOLROYD, The Stephen Perse Foundation, UK

ANJA THIEME, Microsoft Research, UK

HODA HEIDARI, Carnegie Mellon University, USA

DANIELA MASSICETI, Microsoft Research, Australia

JENNIFER WORTMAN VAUGHAN, Microsoft Research, USA

CECILY MORRISON, Microsoft Research, UK

Measurement is essential to improving AI performance and mitigating harms for marginalized groups. As generative AI systems are rapidly deployed across geographies and contexts, AI measurement practices must be designed to support repeatable, automatable application across different models, datasets, and evaluation settings. But the drive to automate measurement can be in tension with the ability for measurement instruments to capture the expertise and perspectives of communities impacted by AI. Recent work advocates for breaking measurement into several key stages: first moving from an abstract concept to be measured into a precise systematized concept; next operationalizing the systematized concept into a concrete measurement instrument; and finally applying the measurement instrument on data to produce measurements. This opens up an opportunity to incorporate community engagement in the systematization phase before operationalizing and applying measurement instruments. In this paper, we explore how to involve communities in systematizing the concept of “cultural appropriateness” in text-to-image models’ representation of culturally significant artifacts through case studies with three communities: blind and low vision individuals residing in the UK, residents of Kerala, and residents of Tamil Nadu. Our systematized concepts reflect community members’ lived experiences interacting with each artifact and how they want their material culture to be depicted. We explore how these systematized concepts can be operationalized into automated measurement instruments that could be applied using a multimodal LLM-as-a-judge approach. We reflect on the benefits and limitations of such approaches.

CCS Concepts: • **Human-centered computing** → **Empirical studies in HCI**.

*Second authors, equal contribution.

[†]Work completed while at Microsoft Research.

Authors’ Contact Information: Nari Johnson, narij@andrew.cmu.edu, Microsoft Research, Cambridge, UK and Carnegie Mellon University, Pittsburgh, PA, USA; Deepthi Sudharsan, deepthi.sudharsan@gmail.com, RiskSpan, Bengaluru, India; Hamna, hamnaabid@gmail.com, Microsoft Research, Bengaluru, India; Samantha Dalal, Sd3765@princeton.edu, Microsoft Research, Cambridge, UK and Princeton University, Princeton, NJ, USA; Theo Holroyd, theo@theoholroyd.com, The Stephen Perse Foundation, Cambridge, UK; Anja Thieme, anthie@microsoft.com, Microsoft Research, Cambridge, UK; Hoda Heidari, hheidari@cmu.edu, Carnegie Mellon University, Pittsburgh, PA, USA; Daniela Massiceti, daniela.massiceti@gmail.com, Microsoft Research, Sydney, Australia; Jennifer Wortman Vaughan, jenn@microsoft.com, Microsoft Research, New York, NY, USA; Cecily Morrison, cecilym@microsoft.com, Microsoft Research, Cambridge, UK.



This work is licensed under a Creative Commons Attribution 4.0 International License.

FACCT '26, Montreal, QC, Canada

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2596-8/2026/06

<https://doi.org/10.1145/3805689.3812222>

Additional Key Words and Phrases: measurement, evaluation, cultural evaluation, multimodal, text-to-image generation

ACM Reference Format:

Nari Johnson, Deepthi Sudharsan, Hamna, Samantha Dalal, Theo Holroyd, Anja Thieme, Hoda Heidari, Daniela Massiceti, Jennifer Wortman Vaughan, and Cecily Morrison. 2026. Evaluating AI-Generated Images of Cultural Artifacts with Community-Informed Rubrics. In *The 2026 ACM Conference on Fairness, Accountability, and Transparency (FAccT '26)*, June 25–28, 2026, Montreal, QC, Canada. ACM, New York, NY, USA, 61 pages. <https://doi.org/10.1145/3805689.3812222>

1 Introduction

People across the world have adopted generative AI tools to automate the creation of images for design and ideation, marketing and advertising, illustration, and beyond [30, 36, 41, 60]. But not all cultures are rendered appropriately. HCI and AI researchers have documented the many ways in which state-of-the-art generative AI systems systematically underperform at depicting historically marginalized cultures, including replicating historical biases in media [41, 74], reinforcing stereotypes [12, 38, 45, 59], and contributing to cultural erasure [40, 61, 76, 93].

Effective model evaluation is necessary to make these failures visible and measure progress toward addressing them [50, 117, 119]. Today, many existing approaches to generative AI evaluation apply off-the-shelf benchmarks and metrics [35, 50, 62]. However, an emerging line of research suggests that these existing approaches to evaluation suffer from significant validity issues [25, 117], and in particular, break down when applied in marginalized, low-data contexts [44, 61, 68]. More generally, these approaches do not attempt to incorporate the expertise or input of the people whose cultures are depicted in the thousands of AI-generated images created each day [74, 93, 94].

An emerging body of human-centered literature has explored ways to involve humans that hold relevant “lived-experience expertise” [78] (e.g., due to their knowledge, standpoint, or cultural identity) in the process of evaluating generative AI. This ranges from efforts that crowdsource culturally specific datasets [70, 76, 88, 96, 100] to those that invite participants to evaluate model outputs [7, 45, 74, 93]. At the same time, there is a rapidly growing literature on ways to develop practical, repeatable, and automatable evaluation practices by using MLLM-as-a-judge evaluations, which delegate evaluative judgments to multimodal language models (MLLMs) in lieu of human judges [101, 110]. This paper aims to connect these two literatures by engaging people with lived-experience expertise in developing rubrics that can be used within the MLLM-as-a-judge paradigm (Figure 1).

Recent work by Wallach et al. [117] calls for AI researchers to reimagine their approach to measurement by adopting an established measurement framework from the social sciences [1]. Social scientists have long grappled with how abstract concepts, like political ideology or job satisfaction, can be captured in precise measurements [1, 26]. Similar questions arise in the evaluation of generative AI systems, where one might wish to measure abstract concepts such as whether a language model can “reason” [98], or whether an AI image depicts a “stereotype” about a social group [12]. The framework breaks the process of measurement into three concrete steps: (1) *systematizing* an abstract concept to be measured into a concrete definition; (2) *operationalizing* the systematized concept into a measurement instrument; and (3) *applying* the measurement instrument on data to produce measurements. The authors point to the opportunity for participants with different forms of expertise to be included in the systematization process, but stop short of offering practical guidance on how to facilitate such engagements. In this work, we ask: how can we scaffold community engagement in the systematization process to develop community-informed measurement instruments?

To answer this question, we facilitate research workshops with members of three communities: blind and low vision individuals residing in the UK and residents of two distinct South Indian states, Kerala and Tamil Nadu. We note that the term “community” has been used in a range of different ways across different academic disciplines [16, 121]. In this work, we follow past HCI studies [11, 48, 78] and use “community” to refer to a group of individuals who hold a shared cultural identity, recognizing the shared “lived-experience expertise that individuals hold that extends beyond their formal training or credentials” [78]. Our goal is to identify ways for non-technical people (people who are not AI developers) to contribute their lived-experience expertise in shaping AI evaluation

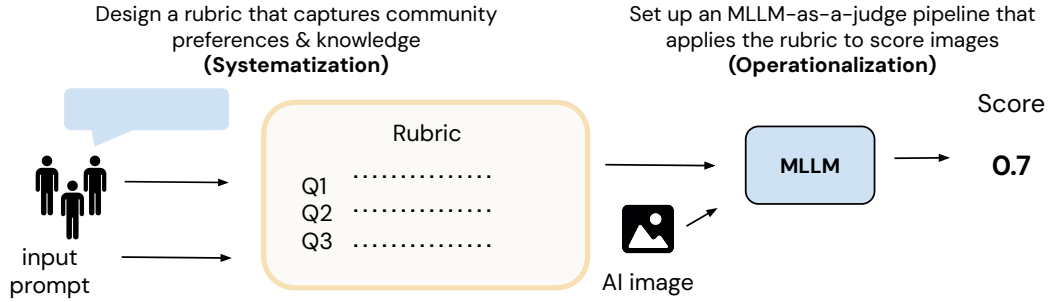


Fig. 1. **Scaffolding community engagement to develop community-centered measures of cultural representation.** Given an input prompt (e.g., “a photo of a guide cane”), we invited community members to participate in designing a rubric that captures their expertise and preferences for each cultural artifact (systematization). Our research team then explored the use of this rubric within an automated multimodal LLM-as-a-judge pipeline (operationalization).

design. In each context, we develop community-informed systematized concepts of *cultural appropriateness* in depicting cultural artifacts [81]. In this work, we view cultural appropriateness loosely as the degree to which AI images of artifacts align with and respect the values, norms, and knowledge systems of the relevant community [93], but leave it to communities to determine what this means. In studying culturally appropriate depictions, our goal is to capture community members’ understandings of how each artifact should—and critically, should not—be represented in AI media [74, 93, 94]. We structure our inquiry around the following research questions:

- **RQ1 (Systematization):** How can practitioners work with community members to systematize the cultural appropriateness of important cultural artifacts? How do the rubrics elicited through a community-centered approach differ from those generated by LLMs?
- **RQ2 (Operationalization):** How might practitioners automate the application of community-informed rubrics to score new AI images?

We contribute an empirical account of how practitioners can engage marginalized communities in the systematization process. We highlight how community members’ expertise, lived and embodied experiences with each artifact, and subjective preferences for representation shape our systematized concepts, demonstrating the potential value of involving communities in measurement. We explore whether rubric application can be automated using an MLLM judge, surface several limitations of MLLM judges, and contribute methodological guidance for how practitioners can assess the feasibility of automating measurement. We conclude by reflecting on open challenges and limitations of our exploratory work, both of which pose opportunities for future research to bring community expertise into the design of AI evaluation metrics.

2 Related Work

Cultural Representation and AI-Generated Media. AI-generated images have the power to shape how communities are understood and perceived by others, as well as how they perceive themselves [46, 94]. Media representation of diverse peoples, especially those who hold historically marginalized identities, is an essential building block for social change and equality [41, 118]. Marginalized communities have used AI tools to create visual media that aligns with their goals for representation [4, 28, 49, 52, 57]. However, these same AI tools have been shown to reproduce normative identities and narratives that contribute to the erasure of already marginalized communities [12, 41, 75]. This can result in potential representational harm through unfavorable or

demeaning depictions that can invoke psychological distress or frustration for users who aim to create media about their own culture [74, 123].

In this paper, we focus on the issue of appropriate cultural representation—an often contested concept [20]—in AI-generated images, as defined by community members. We adopt a broad conceptualization of culture as based on shared identity, recognizing that individuals can inhabit multiple cultural identities simultaneously [133]. Culture may be grounded in place (e.g., shared nationality [2]), or in other identity aspects including race or ethnicity [34], disability [75], sexuality [112], profession [115], or relational roles [116]. Cultural representations—those intended to capture tangible objects (artifacts, monuments) and intangible practices (traditions, rituals) salient to communities’ cultural identity [3, 15, 47]—are not objective constructs with a definitive ground truth. Instead, they involve an articulation of how community members would like their community to be depicted [93]. While community members may have varying perspectives on which aspects of their culture should be highlighted, their knowledge and viewpoint are critical to deciding what makes a representation “appropriate” [45, 93].

In this paper, we focus on designing quantitative measures that reflect participants’ lived understandings of their culture. Prior work in AI and culture has emphasized the need for methodologies that “translate qualitative insights into algorithmic interventions” [13]. Our work responds to this call by bringing qualitative methods into the design of quantitative metrics by engaging communities in the design of rubrics that can be programmatically applied with MLLM-as-a-judge systems.

Measuring the Performance of Generative Text-to-Image Models. Our work contributes to a growing body of interdisciplinary scholarship focused on developing and critiquing performance measures used to evaluate generative AI systems. Drawing from the quantitative social sciences, we define measurement as the assignment of quantitative or qualitative values to specific concepts or properties of a system [98]. In text-to-image generation, concepts that are often measured include generated images’ faithfulness to the user’s generation prompt [56, 99], photorealism [53, 58], and aesthetic value [71]. These measurements are applied to tasks like helping practitioners select deployment-appropriate systems [95, 98, 119] or filtering content below quality thresholds [54, 66, 119].

Despite the recognized importance of evaluation in generative AI development, significant challenges remain. Given the speed at which the AI industry is moving, there is a need for evaluation practices to be repeatable, automatable, and applicable across different models, datasets, and settings [50, 101, 119]. The dominant approach to AI evaluation in industry involves calculating standardized metrics that capture model performance on test datasets—a practice known as “benchmarking” [89, 95, 117]. However, AI benchmarks are often decontextualized and divorced from real-world deployment contexts, leading to validity issues [33, 35, 95]. This decontextualization manifests in the systematic exclusion of content from minoritized communities, including non-English language content and culturally specific imagery [79, 97]. Recognizing these limitations of benchmarks, researchers have called for more localized, disaggregated, contextual evaluations [8, 33]. Many existing efforts expand the scope of generative AI evaluations by curating globally diverse datasets of input prompts [47, 65, 76, 100, 124].

In this work, we shift the focus from datasets to *measures*. It is increasingly common to automate evaluations using model-mediated evaluation approaches, where an auxiliary, pretrained model’s internal representations or judgments are used to stand in for human evaluative judgments as a form of “ground truth” [62]. Drawing on the affordances for scale of LLM-as-a-judge pipelines, practitioners have increasingly turned to multimodal language models (MLLMs) as judges for evaluating generated images [22, 56, 73, 131]. In these pipelines, a multimodal model scores each image using a rubric, which in some cases is generated by another model. Such approaches, however, introduce a fundamental dependency: the quality of the measurements becomes limited by the capabilities and biases of the auxiliary models [62, 68, 76, 77, 100]. Since these models are trained on historical datasets that may lack diversity or contain societal biases, they can systematically undervalue or misrepresent content from marginalized communities. For example, if an auxiliary model was trained primarily on Western

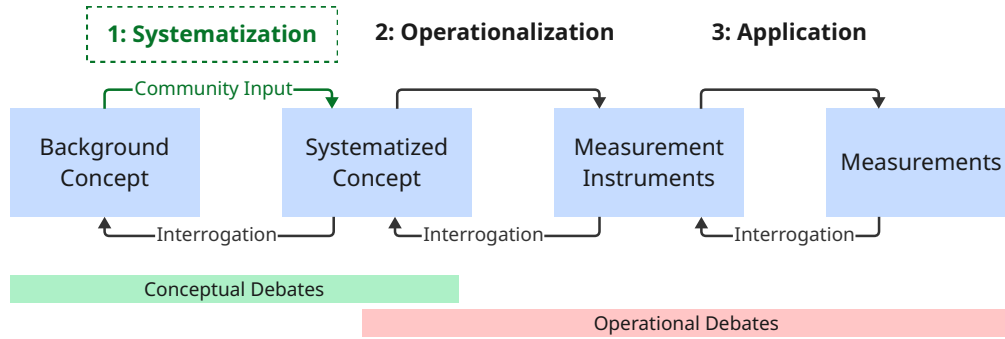


Fig. 2. **Measurement framework from the social sciences** [1, 117]. We study how to center community expertise in the systematization process before operationalizing the systematized concept as an automated MLLM-as-a-judge system.

imagery, it may poorly evaluate AI-generated images of non-Western cultural practices, foods, or aesthetics. We ask whether this limitation might be overcome by the design of more thoughtful rubrics.

Applying Measurement Theory to AI Evaluations. Recent work by Wallach et al. [117] advocates for researchers to rethink how they evaluate generative AI systems, drawing on the tradition of measurement theory from the social sciences. They build on the framework of Adcock and Collier [1], pictured in Figure 2, which shows how abstract concepts can be systematically transformed into concrete, quantifiable measurements. First, during *systematization*, a vague background concept is transformed into a more concrete, well-defined *systematized concept*. Next, during *operationalization*, the systematized concept is developed into a *measurement instrument* (e.g., using metrics or code). Finally, there is *application*, in which the measurement instrument is applied to a set of instances (e.g., outputs of an AI model), resulting in the creation of measurements. This process is intended to be iterative, with iteration and refinement at each stage. As Wallach et al. [117] note, separating systematization from operationalization has the benefit that *conceptual debates* about what is being measured and why can be separated from *operational debates* about *how* a concept will be measured in practice. This allows people with different backgrounds and situated knowledge to be involved in conceptual debates without needing to understand the technical details of how the measurement instrument will be developed, standing in contrast with dominant evaluation practices, where participants are typically only involved at the end of the process, tasked with evaluating the quality of model outputs [56, 61, 90].

While prior work has drawn on disciplinary expertise in systematizing generative AI performance measures—e.g., by consulting psychologists or linguists [24]—there is little work on how to engage community members in translating their lived experience and knowledge into systematized concepts for evaluation. One notable exception is concurrent work by Nguyen et al. [82], who engage community members to iteratively validate and refine a systematized concept of erasure. In this work, we explore how we can scaffold community engagement in the systematization process to develop community-informed systematized concepts that can be operationalized using MLLM-as-a-judge approaches and ultimately applied across different datasets and models.

3 Methods

3.1 Selecting communities and cultural artifacts

We collaborated with members of three distinct communities: blind and low-vision (BLV) people located in the United Kingdom, and current and former residents of two states in South India: Tamil Nadu and Kerala.

We chose these specific communities because members of the research team belong to them. This close researcher–community collaboration supported activities such as artifact selection, participant recruitment, and interpretation of our findings. Studying cultural representation across communities varying in ability, geography, and ethnicity allows us to identify both shared principles and context-specific needs in how communities want to be represented. Recognizing the rich diversity and breadth of experiences that exists within each of these communities, we follow past work [11] to engage community members not as representatives of the entire population with which they share one facet of their cultural identity, but instead as individual experts in their experience as people who hold those identities.

For the BLV community, we focused our study on a single region (the UK). We chose to geographically situate our study because the material culture of blindness is shaped by local factors, such as the public funding of assistive technology use in schools [23]. This broad scoping of eligibility also enabled us to recruit from community organizations with membership bases that were similarly scoped to the UK, such as an email list of UK-based assistive technology users. We also conducted two separate engagements with residents of two Indian states: Tamil Nadu and Kerala. Cultural studies scholars have discussed how India’s 28 region-states function not merely as administrative regions, but also each carry their own unique histories and cultural practices that play an important role in creating a shared cultural identity [102, 111]. For example, residents of Kerala and Tamil Nadu speak different languages (and write in different scripts), and participate in both cultural ceremonies (e.g., festivals like Pongal) and everyday activities (e.g., dance forms like Bharatanatyam) that are unique to their region. We follow past work [100] and center community engagements on cultural artifacts that are *unique to each state*, and do not have a direct equivalent in neighboring geographic states. Each engagement was facilitated by a co-author who was from the respective state.

Guided by the members of our research team who belong to each community, we selected salient cultural artifacts [81]. These artifacts are not intended to represent the richness and breadth of each community’s material culture [133], but to serve as a starting point for our exploration of rubric design. For the BLV community, we selected two assistive technologies: a *guide cane*, a mobility device that helps its user detect obstacles and navigate their surroundings, and a *braille notetaker*, an electronic device used to write and read notes in tactile braille. Prior work has shown that people with disabilities value accurate depictions of assistive technologies in visual media [74, 130] and AI-generated images [75], but text-to-image models fail to accurately depict them [74, 75]. The particular technologies were chosen because of their wide use by and cultural significance to the BLV community in the UK, and are pictured in Figure 3.

For the two Indian communities, we selected objects regarded as meaningful to their state’s cultural identity and likely to be known by residents. From Kerala, we selected (1) the *Kasavu saree*, a traditional white and gold garment, and (2) the *Chundan Vallam*, a snake boat used in traditional boat racing festivals. From Tamil Nadu, we selected (1) the *Pallanguzhi*, a traditional two-player mancala board game, and (2) the *Mridangam*, a percussion instrument that is widely used in South Indian classical music. Prior work has shown failures of text-to-image at depicting both scenes and objects from South Asian cultures [93, 94]. These objects, pictured in Figure 3, carry historic and symbolic value for their roles in each state’s unique cultural ceremonies, festivals, and traditions [81, 125]. We organize workshops by state as residents of neighboring states are less likely to be familiar with these objects. For instance, the Chundan Vallam is unique to the Keralan water festival tradition, and would not be widely known by residents of neighboring Tamil Nadu.

3.2 Eliciting culturally appropriate artifact depictions in workshops

To create the rubrics, we conducted a series of synchronous, hour-long online workshops with community members to elicit their knowledge of and desires for representation of the selected cultural artifacts. Workshops were facilitated by members of our research team, and took place between December 2024 and April 2025. To recruit participants, we used a purposive sampling approach [91], drawing on the research team’s existing networks



Fig. 3. **Selected culturally significant artifacts.** From right to left: (1) With the blind and low vision community, we selected a *guide cane* (a mobility aid that is held diagonally across one’s body) and a *braille notetaker* (an electronic device that can be used to read and write notes in tactile braille). (2) With residents of Tamil Nadu, we selected *Pallanguzhi* (a two-player mancala game where players compete to collect cowry shells or seeds) and *Mridangam* (a percussion instrument widely used in South Asian classical music). (3) With residents of Kerala, we selected a *Kasavu saree* (a handwoven saree from Kerala, known for its off-white body with a gold border), and *Chundan Vallam* (a traditional boat from Kerala with a raised prow commonly used in festival races).

across each community. Our final sample included 10 BLV participants in the UK, 9 participants from Tamil Nadu, and 8 participants from Kerala. Participants were compensated either £75 (UK) or 500 INR (India) based on their location, and all interviews were conducted in English. All workshop studies were approved by our institution’s ethics review board. We detail the structure of the workshops, including activity guides, participant demographics, and steps we took to ensure the workshops were accessible for members of the BLV community in Appendix B.2.1.

For each artifact, we curated a small set of AI-generated images and real photographs to facilitate discussions of cultural representation. We generated images using two state-of-the-art text-to-image models at the time of the study, Stable Diffusion 3 [105] and DALL·E 3 [84], prompting each model with a simple template (e.g., “A photo of a {artifact name}”). For each artifact, we generated an initial pool of images and then grouped images by shared visual characteristics, sampling across groups to create a manageable but meaningfully diverse set for participants to review [74]. In addition, we curated real photographs of each artifact from trusted sources (e.g., the Royal National Institute of Blind People [113]). These photographs helped ground discussion by providing clear examples of culturally appropriate depictions. For instance, we could more easily locate a photograph of a braille notetaker that depicted a QWERTY input keyboard correctly than generate one. We could also curate photographs that sat near the boundary of what participants might consider acceptable to facilitate discussion. The final set of images used in the workshops is included in Appendix B.

For each artifact, we conducted two study activities. First, after presenting each image (by screen-sharing and/or presenting an alt-text description for BLV participants), the facilitator asked participants to share if they felt that the image could or could not be shown to the general public to represent the artifact and why. The facilitator then prompted each participant to elaborate on what made a particular image a good, bad, offensive, or incorrect depiction of the artifact. We discussed 10 total images of each assistive technology, and 16 total images of each Indian artifact.

Second, we invited participants to reflect on what they had seen so far to create a list of the most important criteria for a culturally appropriate portrayal of the object and provide reasons for each of their responses. This activity encouraged participants to articulate concrete visual criteria that shaped their decisions. The study

facilitator encouraged participants to reflect on the degree of acceptable variability for objects and how visual features should be prioritized.

3.3 Systematizing cultural representation into rubrics

We used the data collected from workshops to develop systematized concepts of “cultural appropriateness” for each artifact. These took the form of simple rubrics made up of binary yes/no criteria that community members identified as integral to an appropriate representation of the artifact; see Figure 4 for an example of the rubric produced for a guide cane. We adopted this structure because rubrics are a common measurement approach (e.g., in domains beyond AI, such as the social sciences [1]), and they are compatible with existing MLLM-as-a-judge pipelines [22].

To distill criteria for the rubrics, we first organized participants’ statements from both activities into two categories: (1) *criteria*, or the concrete visual features or physical characteristics that shaped participants’ decision-making, and (2) *justifications* that participants provided for their decisions. In deciding which criteria to include in the rubrics, we made two key decisions. First, we excluded criteria that were highly contested across participants. For instance, community members disagreed about whether a guide cane with red reflective tape on its body was an appropriate representation. As such, we did not include a criterion about red tape in our rubric. We also excluded criteria that participants consistently described as less essential to depicting the essence of each artifact. Consequently, rather than offering a complete description of each artifact, our rubrics focus on the 5 to 10 core features that participants agreed were most essential to a culturally appropriate depiction. We acknowledge that this choice is value-laden, and discuss alternative ways that practitioners might navigate and resolve contested criteria in Section 5.3. Importantly, the resulting rubric criteria for each cultural artifact should be interpreted as being *particular to* (rather than portable across) the specific cultural contexts that they are created in collaboration with.¹

Finally, we grouped the criteria into higher-level themes around the motivation for their inclusion. These themes reflect community-specific understandings of what constitutes (in)appropriate cultural representation. For example, BLV community members emphasized that assistive technologies should be usable and accessible to blind users, clarifying why particular visual features were prioritized. We structured each rubric by these themes.

Using the rubric, we say that an image is “culturally appropriate,” receiving a label of 1, only if it satisfies all criteria. Otherwise, we say that an image is culturally inappropriate (label 0). While these final binary cultural appropriateness labels do not capture the subjective and contested nature of cultural representation, we adopt this aggregation function for its simplicity to calculate, and its direct comparability with participants’ judgments.

As a point of comparison, we also generated rubrics automatically, in line with common AI industry practices [22, 101, 109]. Specifically, we prompted GPT-4o to produce evaluation criteria describing the “most important visual characteristics that should be present or absent” in culturally appropriate depictions of each artifact (details in Appendix C.2). To compare our community-informed rubrics with the automatically generated rubrics, the first author worked with community members on the research team to annotate each rubric, with the aim of identifying overlaps, divergences, and omissions relative to both our community-informed rubrics and our own knowledge of the artifacts. We then assessed each criterion for its alignment with community perspectives, and identified broader patterns in what LLM-generated rubrics captured or overlooked.

3.4 Operationalizing the rubrics through MLLM-as-a-judge

To operationalize our systematized concepts, we explore whether our community-elicited rubrics can be automatically applied using an MLLM-as-a-judge pipeline. We adopt a simple system prompt, following past work [22, 132], that instructs the MLLM to “*determine whether the provided image meets each criterion.*” We adopt GPT

¹For instance, a guide cane rubric created based on engagements facilitated in a different part of the world (e.g., with residents of Tanzania) may result in a different set of scoring criteria that reflect a different set of participant experiences and desires for cultural representation [45].

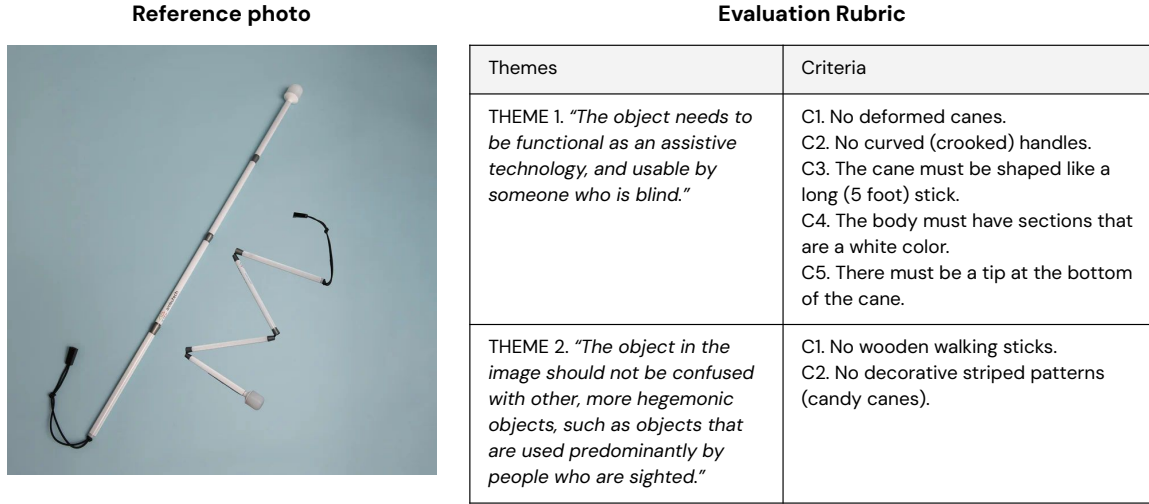


Fig. 4. **A rubric to score images of a guide cane, designed with BLV community members.** Criteria that correspond to visual features in images are organized under two themes that describe participants’ desires for cultural representation.

4-o [86] as our judge model for its demonstrated performance on MLLM-as-a-judge tasks [131] and report all results averaged over five random seeds. We provide additional details in Appendix C.3.3.

To validate the operationalization—that is, whether the MLLM-as-a-judge measurement instrument captures the systematized concept—we compare judgments of the rubric criteria from an MLLM judge and human annotators on a new dataset of AI-generated images. We create a dataset containing 50 images of each artifact generated from 5 different models (DALL-E 3, Stable Diffusion 3 Medium, Stable Diffusion 3.5 Large, GPT Image-1, and Flux.1 Dev) using simple prompts (like “a photo of a guide cane”). We include image generation details in Appendix C.3.1. These images depict the types of representations and errors that state-of-the-art models produce today. For each image, a member of our research team manually annotated whether each rubric criterion is met, which we compared against annotations assigned by the MLLM judge. Each rubric was applied by the research team member who facilitated workshop engagements and led the creation of the rubric, to ensure its consistent application across the dataset. We include an extended discussion of manual (human) application of our rubrics and inter-annotator agreement in Appendix C.4.

4 Results

In this section, we present our community-informed rubrics (Section 4.1) and an exploration of the feasibility of applying these rubrics automatically using an MLLM-as-a-judge approach (Section 4.2). We reflect both on key trends that are common across the three communities we engaged with, and on key differences across communities and artifacts.

4.1 Systematizing community expertise into evaluation rubrics (RQ1)

In this section, we present the rubrics used to systematize the concept of “culturally appropriate” representation for each cultural artifact. The rubric for a guide cane is shown in Figure 4; the remainder can be found in Appendix A. For each community, we first identified high-level themes that capture key dimensions of how participants

evaluated cultural representation. We used these themes to organize a set of criteria that correspond to concrete, observable visual features.

4.1.1 Themes and criteria to assess cultural representation. Across all three communities, participants consistently evaluated images along two core dimensions: *functionality*, whether an artifact could plausibly serve its intended purpose, and *recognizability*, whether the depiction preserves the key features that distinguish the artifact from related, but culturally distinct objects. Within the BLV context, participants drew on their embodied experiences as users of each assistive technology to assess whether a depiction would be able to serve its intended function, e.g., by imagining how one might hold a pictured cane or whether depictions of braille could be read by touch. Participants repeatedly pointed out when generated images “looked like” other recognizable objects that did not serve the intended purpose of the assistive technology, such as a generated image of a guide cane that resembled a walking stick, a distinct mobility aid that differs in its function, material, and handle shape. When asked to elaborate, participants explained that a walking stick is predominantly used by people who are sighted. Community members felt that these “confused” depictions were not only misleading, but were disappointing in that they failed to communicate the unique material culture of blindness.

We found many key similarities in the higher-level themes that describe how residents of Kerala and Tamil Nadu evaluated their cultural artifacts. Participants from both Kerala and Tamil Nadu emphasized evaluating each artifact’s functionality by assessing whether the artifact pictured could be used for its intended purpose. For instance, participants from Tamil Nadu noted how the Mridangam drum’s leather straps running along its body are critical for tuning the instrument when played. Similarly, participants from Kerala described the unique shape of the Chundan Vallam racing boat (e.g., its narrow, pointed stern) as central to its performance, and agreed that depictions should include the Amarakaran: the standing oarsmen who “orchestrates and steers the team’s focus during the race”. Similarly, participants from both Indian states repeatedly pointed out when generated images failed to capture what they felt were the most culturally distinct features of each artifact, such as the golden hand-woven border that distinguishes a Kasavu saree from sarees emblematic of other regions (“if it’s not white with a gold border, then it’s just a saree”), or confused depictions that resembled other board games (like checkers) instead of Pallanguzhi.

Taken together, these two shared dimensions of functionality and recognizability enable us to unpack what is shared across all three communities in their desires for cultural representation, while including rubric themes and criteria that remain context-specific and grounded within the unique perspectives of each community.

4.1.2 Assessing the alignment of our rubrics with community preferences. We interrogate the validity of our rubrics by examining whether they align with preferences expressed by community members during workshops. Specifically, we compare (a) labels of cultural appropriateness obtained by having a member of our research team manually evaluate the rubric criteria to (b) participants’ judgments of whether or not images were appropriate to be shown from the first workshop activity. We aggregate judgments across workshop participants by taking the majority judgment.² We find that for over 80% of images, the label produced by manually applying the rubric matches the majority participant judgment (see Appendix Table 8). Note that we may not expect perfect alignment because of the inherently plural and contested nature of cultural representation, which lacks a singular ground truth. In our own workshops, participants’ judgments of cultural appropriateness disagreed with each other for 33% of the workshop images.

4.1.3 Comparing our community-informed rubrics with LLM-generated rubrics. We observe several differences when examining how our rubrics differ from those generated by an off-the-shelf LLM, as described in Section 3.3. Many of the LLM-generated rubric criteria accurately capture surface-level properties of the artifacts, such

²We note as one limitation of this comparison that these participant judgments were used to create the rubrics in the first place. Ideally we would interrogate validity using a new set of images and new participant judgments.

Table 1. **Human vs. MLLM application of the rubrics.** Using 50 generated images per artifact, we report (i) the proportion of images labeled as culturally appropriate by humans versus an MLLM, and (ii) agreement between MLLM and human labels. We further disaggregate agreement by images the human labeled as appropriate versus inappropriate.

Artifact	Human	MLLM	Agreement	Agreement	Agreement
	(% Appropriate)	(% Appropriate)	Overall	Appropriate images	Inappropriate images
Guide cane	0.40	0.44	0.84	0.84	0.83
Braille notetaker	0.08	0.20	0.82	0.65	0.83
Pallanguzhi	0.18	0.12	0.78	0.22	0.90
Mridangam	0.10	0.21	0.84	0.76	0.85
Kasavu saree	0.12	0.21	0.88	0.87	0.88
Chundan Vallam	0.00	0.17	0.83	N/A	0.83

as noting the “distinct golden border” of a Kasavu saree or describing a guide cane as “a long, slender stick” (Appendix C.2). However, several LLM-generated rubrics include criteria that are inaccurate, reflecting a fundamental misunderstanding of the objects. For example, the rubric for braille notetaker requires that the depiction “resemble an electronic notebook” and “include a display for visual feedback,” two features that are inaccessible to blind users and uncharacteristic of braille notetakers. More generally, most LLM-generated rubrics omitted or underspecified features that community members viewed as essential for functionality or recognizability. The Mridangam rubric, for example, omitted the drum’s characteristic black circular membrane that is necessary to its tone. These omissions were not uniform across artifacts: some LLM-generated rubrics, such as that for the Kasavu saree, overlapped substantially with community rubrics, whereas others, most notably the braille notetaker, diverged sharply (see Appendix C.2 for further discussion). This suggests that community participation is essential for artifacts or concepts that are at present poorly captured by LLMs.

4.2 Operationalizing MLLM-as-a-judge metrics (RQ2)

We demonstrate how our community-informed rubrics could be operationalized into automated measurement instruments by using an MLLM to assess the rubric criteria, as described in Section 3.4. We also consider a manual implementation, in which criteria are annotated by a member of our research team. We use both of these measurement instruments (automated and manual) to label 300 images generated by five state-of-the-art text-to-image models as culturally appropriate or not. Table 1 summarizes the results. In what follows, we first analyze the human annotations of criteria to examine how state-of-the-art models depict cultural artifacts, showing how the rubrics enable interpretable comparisons that reveal model-specific failure modes. We then compare MLLM-generated and human annotations. We emphasize that our goal in this work is not to optimize this implementation of the MLLM-as-a-judge system, but rather explore its feasibility, as best practices for MLLM-as-a-judge approaches remain an active area of research [21, 72, 101].

4.2.1 What rubrics reveal about cultural representation. We begin by examining the human annotations of the rubric criteria for the 50 images per artifact. Across all artifacts, we find that state-of-the-art text-to-image models rarely produce depictions that satisfy all rubric criteria, resulting in low rates of images being labeled as culturally appropriate (first column in Table 1). The proportion of culturally appropriate images varies substantially across artifacts—40% of guide cane images meet all criteria, compared to roughly 10% of braille notetaker and Mridangam images, and no images of the Chundan Vallam. This pronounced class imbalance for all but one artifact reflects the current limitations of frontier models in representing cultural artifacts. This indicates that our community-informed rubrics capture meaningful errors of representation, even for models that we did not show to participants during our workshops, such as GPT Image-1 and Flux (discussed further in Appendix C.5).

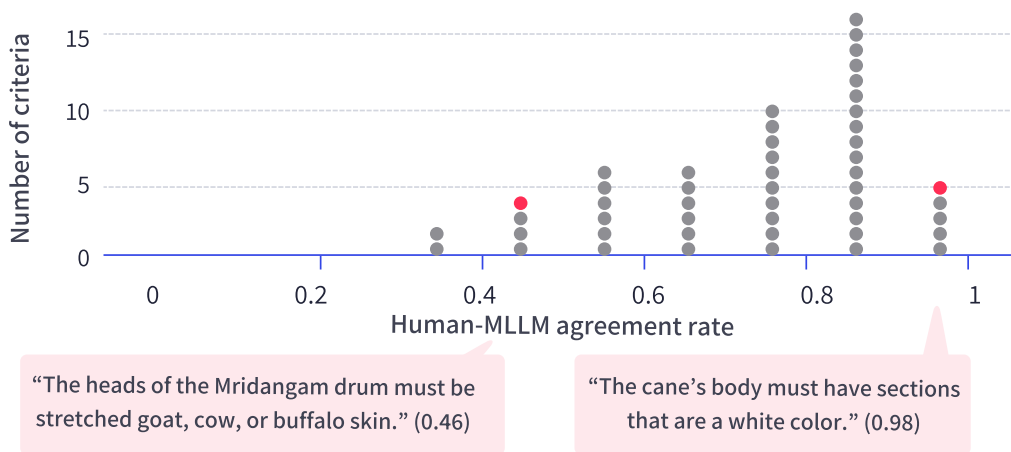


Fig. 5. **Human-MLLM judge alignment for individual rubric criteria.** A histogram that shows the human-MLLM agreement rate for individual rubric criteria. We find that there is high variance in the MLLM’s ability to annotate criteria accurately. For instance, GPT 4-o has low accuracy (agreement rate 0.46) annotating whether a drum’s head is made of the correct material, but high accuracy (agreement rate 0.98) determining whether a cane is white in color.

Inspecting the specific criteria that are commonly violated, we find that different models exhibit distinct failure modes (see Section C.5 in the appendix). For example, the GPT Image-1 images of the Mridangam are primarily deemed inappropriate because they (incorrectly) add atypical decorative patterns on the drum’s body. In contrast, the other models violate a much wider range of rubric criteria: e.g., applying our rubrics reveals how all of the DALL-E 3, Flux.1 DEV, and Stable Diffusion images fail to depict the Mridangam’s black circular membrane. This analysis highlights the potential value of our rubric criteria in providing interpretable insights about the types of errors of representation made by different models, and the ability to track and measure them over time.

4.2.2 Interrogating the validity of automated application. To assess the feasibility of automating rubric application, we compare MLLM-generated annotations to human annotations for each rubric criterion. Across artifacts, the MLLM’s labels of cultural appropriateness agree with human judgments for 78–88% of the images in our dataset (fourth column in Table 1), with a consistent tendency to over-predict appropriateness for five of the six artifacts. Disaggregating by images labeled as appropriate versus inappropriate by humans reveals two types of MLLM errors. Human-MLLM agreement for the *inappropriate* images tells us the MLLM’s ability to recognize violated criteria. In contrast, human-MLLM agreement for the *appropriate* images tells us whether the MLLM can recognize valid depictions. Each has different implications for how practitioners might revise their rubric criteria to improve judge performance, as the best steps forward depend on the type of error being made—for instance, relaxing or clarifying criteria to address false negatives vs. strengthening or expanding criteria to address false positives.

To better understand the differences between the manual and automated implementations, we break down agreement by each of the individual rubric criteria. Figure 5 shows the distribution of agreement scores across the 48 total criteria taken from our six rubrics. Agreement varies substantially: while many criteria are annotated accurately (44% of criteria have agreement rates above 80%), 10% of criteria have agreement rates below 50%, indicating performance no better or even worse than chance. Many of the criteria on which the MLLM-judge performed poorly (i.e., the “low-accuracy criteria”) were those that required the artifact to be depicted in a specific

shape or spatial arrangement, such as a valid braille cell configuration (39% agreement), the stern shape of a Chundan Vallam (43%), or arrangement of pits on a Pallanguzhi board (64%). Other low-accuracy criteria are those based on features that are difficult (for both humans and machines) to infer from visual information alone [37], such as the material of a drum’s head (46%), or the type of wood used to create a Pallanguzhi board (56%). In contrast, high-accuracy criteria typically describe visually salient features, such as the color of a guide cane (98%). Many low-accuracy criteria correspond to features that community members identify as highly important, highlighting opportunities to improve MLLM-based judging. We include the agreement rate for each criterion and a further discussion of judge performance in Appendix C.6.

5 Discussion

In this section, we discuss our study’s contributions to and implications for the emerging interdisciplinary debate around how to bring stakeholder perspectives into the design of evaluations for generative AI, a topic that is gaining attention within the FAccT community [63, 82]. Our goal in conducting this study is not to advocate for any one particular methodology as the ideal way for practitioners to engage stakeholders in measurement design. Instead, our goals are to highlight the key benefits and challenges that surfaced across engagements with three different communities, to offer reflections on the methodological limitations of our exploratory approach, and to identify opportunities for future work. To this end, we reflect on open research challenges in automating the application of rubrics (Section 5.1), situate our work within the larger discourse on tensions between participation and scale (Section 5.2), and discuss important limitations of our study (Section 5.3).

5.1 Towards the design of valid, community-centered MLLM-as-a-judge pipelines

A gold-standard approach to centering community participation in measurement would involve sustained community engagement across all phases of the measurement lifecycle (Figure 2), but we recognize that such approaches might be prohibitively costly for both communities and industry practitioners [31, 128]. Accordingly, our methodology reflects a pragmatic design choice: we invite community members to participate in systematization through one-time workshops, and then examine the extent to which MLLM judges can operationalize these criteria effectively off-the-shelf, without further customization or experimentation. Our findings, however, reveal substantial variation in how reliably MLLM judges apply different rubric criteria, raising important concerns about their resulting validity. Prior work has documented the brittleness of MLLM-based judges, including systematic biases [73, 126] and sensitivity to prompt formulation [72]. More broadly, both AI researchers and quantitative social scientists emphasize that the measurement design process is inherently iterative [1, 39, 101, 117]. As Adcock and Collier [1] argue, systematized concepts often require revision once confronted with the practical challenges of operationalization, e.g., the realization that a systematized concept may be too difficult to reliably measure.

Iteratively refining rubric criteria to improve their legibility to automated judges is a critical next step for future work and a limitation of this current study that aimed to focus on systematization. Addressing these questions requires grappling with both technical challenges in the design of MLLM-as-a-judge pipelines [39, 72], and human-centered challenges in thoughtfully engaging community members in a highly specialized and technical measurement design process [107]. Future work can also conduct more critical studies of the implicit limitations of automated approaches to measurement, e.g., drawing upon literature from critical algorithm studies [64, 120], and provide guidance for when alternative ways of evaluating AI systems (e.g., direct engagement with stakeholders [74, 119]) may be preferred.

5.2 Defining the scope and purpose of community-centered measurement

Our goal in this work is to scaffold community participation in the evaluation of general-purpose, universally-scoped foundation models. Our methodology builds upon the framework put forward by Suresh et al. [108],

who pose that even though foundation model design is largely centralized, local communities can still pursue meaningful participatory efforts at the “surface layer,” which corresponds to a specific context (e.g., generating culturally representative media that depicts the unique material culture of Kerala) and groups of stakeholders (e.g., residents of Kerala whose culture is being depicted in those images).

Our findings show that community participation can shape measurement design by contributing lived, embodied expertise that is difficult to capture through reference images or automated approaches alone. Rather than providing a surface-level description of the visual attributes of each artifact, participants’ descriptions spoke to how artifacts are used and experienced in practice. This insight affirms past work showing that anonymous crowd workers often lack the cultural expertise required for evaluative tasks, even when provided with reference images [45, 47, 93]. In contrast, many LLM-generated rubrics failed to capture these salient features and in some cases reflected confused or inaccurate understandings. This finding extends those from past work, which describe how LLM-generated rubrics are often overly vague or underspecified [109], to reveal how these rubrics may reflect deeper cultural misalignment. Taken together, these results suggest that community participation in defining evaluation criteria can help bridge epistemic gaps in large pretrained models, particularly for low-data cultural artifacts [47, 77], improving the validity of the resulting measurement instruments.

Beyond asking community members to contribute their knowledge, community participation also enables communities to express their more normative desires and preferences for cultural representation. For example, while it is common for braille notetaker devices to have either a QWERTY or braille keyboard, blind community members consistently expressed a preference for braille keyboards, which better emphasize tactile culture and align with shared advocacy goals around increased public investment in accessible technologies. Thus, even in settings where multiple depictions may be technically accurate, participation enables communities to make critical decisions about how they desire to be represented. This more normative role reveals how participation is still valuable in a setting where an LLM can fully describe a cultural artifact, as it is still up to the community to make choices about which depictions they do, or do not, desire.

While our findings illustrate the potential benefits of human-centered approaches to measurement through close collaborations with three communities and just a small set of artifacts, many open questions remain. In real-world settings, foundation model providers must prioritize engagement across thousands of communities and a wide range of content [119, 127, 128], a scale at which both our methods and related qualitative approaches to identifying representational harms (e.g., [74, 93]) are not feasible to apply. Our findings highlight key considerations for identifying where community participation is most valuable in measurement design. We found that LLM-generated rubrics may serve as a reasonable starting point for some cultural artifacts (Section 4.1.3), but not others. Practitioners may consider inviting community members to assess the face validity of LLM-generated rubrics as a lightweight check. If the rubrics diverge significantly from community members’ understandings of the content being evaluated, this is a signal that a deeper engagement (using methods like ours) may be warranted.

In this study, systematization of knowledge into rubrics was done synchronously, and in detail. The depth of our exploratory engagement reveals several promising directions for adapting this process to reduce the potential time and labor expected of both communities and practitioners. One direction is to elicit community members’ desires for cultural representation asynchronously, drawing upon known deliberation approaches (e.g., the Delphi technique [55]) or emerging interfaces to support users in designing LLM-as-a-judge pipelines [39, 92, 101]. Practitioners can also avoid starting from scratch by instead asking communities to revise and critique LLM-generated rubrics, treating automated outputs as provisional baselines rather than authoritative representations. Another complementary direction for future work can explore how practitioners can ground rubric development in themes identified by prior qualitative research, e.g., shared themes that articulate participants’ desires for cultural representation, like those identified here and by Qadri et al. [93], or work that documents the representational harms experienced by marginalized communities [10, 12, 75, 76, 94]. Practitioners can use these themes identified by past literature to scaffold more efficient elicitation of specific visual criteria that can be used to score images.

While such lower-touch approaches may improve scalability, our study suggests that synchronous engagement plays an important role in fostering participant buy-in and surfacing deeper insights into participants' more normative desires for cultural representation (the higher-level themes) [93], highlighting an open question about how much of this depth can be preserved in more lightweight models of participation.

5.3 Limitations

There are many approaches to scaffolding engagements with impacted communities—for instance, pursuing sustained engagements with organizations with defined memberships and structures of communication [114, 128]. In contrast, in this work, we chose to engage individuals who held shared identities and life experiences but otherwise did not know each other. While this recruitment strategy is in line with past HCI research on representational harm [11, 75, 93], our broad conceptualization of “community” as a group of individuals who hold a shared cultural identity has limitations. Our workshop data allowed us to identify preferences that were shared across the individuals we interviewed. However, our small sample of participants recruited through convenience sampling does not enable us to make claims that will generalize across the entire community: a common limitation of qualitative research [69].

We made the value-laden choice to scaffold participation as consultation [6]—one-time workshops to elicit participants' preferences—rather than engaging the community as full collaborators with the power to shape and own research outputs [31, 108]. While there are several potential benefits to pursuing lower-touch approaches to participation, such as respecting and attempting to minimize the labor required from participants [128], we acknowledge that such approaches run the risk of being exploitative if outputs are misused by researchers, or if community members are not adequately compensated for their labor [27, 31]. There is more work needed to develop methods that meaningfully shift power to community members, e.g., as part of a grassroots, community-led project where community members have full ownership over critical measurement decisions [88].

Our decision to prioritize criteria that were agreed upon across participants when designing rubrics, excluding those that were contested, has several limitations. While our approach focuses on those criteria viewed as most essential to a culturally appropriate depiction, without engaging more deeply with contestedness and disagreement, it runs the risk of replicating dominant or hegemonic views [32]. Future research can build upon our methods, which *surface* disagreement across participants, to pursue alternative approaches to *reconcile* these disagreements through continued engagements, e.g., using deliberative or discursive methods to reach group consensus [11, 93], or exploring innovative approaches to measurement that allow for more than one singular “ground truth” [42, 104]. Finally, our rubrics systematize cultural appropriateness for depictions of cultural artifacts in isolation, and further work is needed to extend these methods to design rubrics for more realistic, complex cultural scenes [10, 93, 114].

6 Conclusion

Rich qualitative work from the FAccT community has demonstrated that AI models fail to represent people from marginalized communities as they wish to be seen. We explore one potential path forward: supporting participation in designing evaluations that can be automatically applied to score AI-generated outputs. Our findings open a broader discussion about the technical challenges of designing valid measures, and the value community input adds relative to fully automated evaluation approaches. We are hopeful that practitioners can adapt and extend the exploratory methods presented in this work to create more contextually grounded AI evaluations.

7 Endmatter Statements

7.1 Ethical Considerations Statement

Our study protocol, recruitment material, and consent form were reviewed and approved by our institution's ethics review board to ensure they followed best practices. Participants were compensated and informed that they could withdraw from the study at any time without consequence.

A central ethical consideration throughout the study was co-creating access with participants with differing needs. Drawing on best practices in accessibility research [9, 28, 80], we adapted our protocol to support non-visual access to AI images by creating carefully designed alt text, and inviting BLV community members to participate in cross-ability pairs, detailed in Appendix B.2.1. We also regularly solicited participant feedback during workshops so that we could modify our facilitation approach to increase comfort for participants, e.g., by repeating our alt text descriptions or offering more regular breaks to rest in between activities.

When designing our protocol, we carefully considered the risks of exposing participants to harmful or offensive AI-generated depictions, as prior work has shown that repeated exposure to representational harms can be disempowering [122, 129]. To mitigate these risks, we pre-generated all images shown in workshops rather than generating them live. Pre-generation also allowed our research team to curate a diverse set of meaningfully distinct images that we could show participants, reducing the annotation burden expected from participants so that they would not need to give us feedback on the same types of errors.

Our approach inherits familiar risks and limitations of related participatory AI efforts that are discussed in past scholarship [31, 103, 108, 128]. First, we our proposed methods require labor from participants who may already be marginalized. To this end, we encourage practitioners to consider how participation can be meaningfully compensated, and discuss ways in which participation might otherwise be made less burdensome. While participation in measurement offers great potential to shape foundation model development, we also caution that participatory evaluation may not yield immediate or tangible benefits for community members, such as improved representations in deployed models [27, 103]. Accordingly, we urge practitioners to be transparent about how the outputs of participation (e.g., the resulting measurement instruments) will be used and what benefits, if any, participants can reasonably expect (e.g., beyond contributing to research).

7.2 Generative AI Usage Statement

Generative AI tools were not used to generate original content, analyses, arguments, findings, or interpretations. All academic contributions, including our literature review, workshop data collection, workshop data analysis, interpretation, and writing were produced by the authors. The authors did make use of generative AI tools to support brainstorming throughout the paper writing process. ChatGPT (based on GPT-5.2) was used to brainstorm alternative terminology choices, for instance, for describing measurement approaches that can be deployed at scale. ChatGPT was also used to brainstorm ways to more clearly format results tables and figures. For all uses, the authors did not copy the text directly output by LLMs, but instead reviewed what was generated and used their own judgment to make revisions to the paper.

7.3 Acknowledgements

We thank our study participants from each community who shared their knowledge and expertise. We thank Hanna Wallach, Siobhan Mackenzie Hall, Xinnuo Xu, Melanie Pradier, Martin Grayson, Camilla Longden, Rita Marques, Shaily Bhatt, Shivani Kapania, Luke Guerdan, Fernando Diaz, and Anna Kawakami for helpful feedback and conversations that shaped this work. NJ and HH acknowledge support from the NSF (IIS2229881) and the CMU Block Center for Technology and Society. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not reflect the views of the National Science Foundation and other funding agencies.

References

- [1] Robert Adcock and David Collier. 2001. Measurement validity: A shared standard for qualitative and quantitative research. *American Political Science Review* 95, 3 (2001), 529–546.
- [2] Muhammad Farid Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Singh, Alham Fikri Aji, Jacki O'Neill, Ashutosh Modi, and Monojit Choudhury. 2024. Towards Measuring and Modeling "Culture" in LLMs: A Survey. <https://arxiv.org/abs/2403.15412>
- [3] Muhammad Farid Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Singh, Ashutosh Dwivedi, Alham Fikri Aji, Jacki O'Neill, Ashutosh Modi, and Monojit Choudhury. 2024. Towards Measuring and Modeling "Culture" in LLMs: A Survey. *arXiv preprint arXiv:2403.15412* (2024).
- [4] Rudaiba Adnin and Maitraye Das. 2024. "I look at it as the king of knowledge": How Blind People Use and Understand Generative AI Tools. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility* (St. John's, NL, Canada) (ASSETS '24). Association for Computing Machinery, New York, NY, USA, Article 64, 14 pages. doi:10.1145/3663548.3675631
- [5] Afra FeYZa Akyürek, Advait Gosai, Chen Bo Calvin Zhang, Vipul Gupta, Jaehwan Jeong, Anisha Gunjal, Tahseen Rabbani, Maria Mazzone, David Randolph, Mohammad Mahmoudi Meymand, Gurshaan Chattha, Paula Rodriguez, Diego Mares, Pavit Singh, Michael Liu, Subodh Chawla, Pete Cline, Lucy Ogaz, Ernesto Hernandez, Zihao Wang, Pavi Bhattar, Marcos Ayestaran, Bing Liu, and Yunzhong He. 2025. PRBench: Large-Scale Expert Rubrics for Evaluating High-Stakes Professional Reasoning. <https://arxiv.org/abs/2511.11562>
- [6] Sherry R. Arnstein. 1969. A Ladder of Citizen Participation. *Journal of the American Institute of Planners* 35, 4 (1969), 216–224.
- [7] Lora Aroyo, Alex S. Taylor, Mark Díaz, Christopher M. Homan, Alicia Parrish, Greg Serapio-García, Vinodkumar Prabhakaran, and Ding Wang. 2023. DICES dataset: diversity in conversational AI evaluation for safety. In *Proceedings of the 37th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) (NIPS '23). Curran Associates Inc., Red Hook, NY, USA, Article 2321, 13 pages.
- [8] Solon Barocas, Anhong Guo, Ece Kamar, Jacquelyn Kronen, Meredith Ringel Morris, Jennifer Wortman Vaughan, W Duncan Wadsworth, and Hanna Wallach. 2021. Designing disaggregated evaluations of ai systems: Choices, considerations, and tradeoffs. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 368–378.
- [9] Cynthia L. Bennett, Erin Brady, and Stacy M. Branham. 2018. Interdependence as a Frame for Assistive Technology Research and Design. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility* (Galway, Ireland) (ASSETS '18). Association for Computing Machinery, New York, NY, USA, 161–173. doi:10.1145/3234695.3236348
- [10] Cynthia L. Bennett, Shaun K. Kane, and Christina N. Harrington. 2025. Toward Community-Led Evaluations of Text-to-Image AI Representations of Disability, Health, and Accessibility. In *Proceedings of the 5th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization* (EAAMO '25). Association for Computing Machinery, New York, NY, USA, 256–270. doi:10.1145/3757887.3763012
- [11] Stevie Bergman, Nahema Marchal, John Mellor, Shakir Mohamed, Iason Gabriel, and William Isaac. 2024. STELA: a community-centred approach to norm elicitation for AI alignment. *Scientific Reports* 14, 1 (2024), 6616.
- [12] Federico Bianchi, Pratyusha Kalluri, Esin Durmus, Faisal Ladhak, Myra Cheng, Debora Nozza, Tatsunori Hashimoto, Dan Jurafsky, James Zou, and Aylin Caliskan. 2023. Easily Accessible Text-to-Image Generation Amplifies Demographic Stereotypes at Large Scale. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 1493–1504. doi:10.1145/3593013.3594095
- [13] Asia Biega, Georgina Born, Fernando Diaz, Mary L. Gray, and Rida Qadri. 2025. Towards a Multidisciplinary Vision for Culturally Inclusive Generative AI (Dagstuhl Seminar 25022). *Dagstuhl Reports* 15, 1 (2025), 33–49. doi:10.4230/DagRep.15.1.33
- [14] Black Forest Labs. 2024. FLUX. <https://github.com/black-forest-labs/flux>.
- [15] Janet Blake. 2000. On Defining the Cultural Heritage. *International & Comparative Law Quarterly* 49, 1 (2000), 61–85.
- [16] Emory S. Bogardus. 1942. *Fundamentals of Social Psychology* (3 ed.). D. Appleton-Century Company, New York and London.
- [17] Stacy M. Branham and Shaun K. Kane. 2015. The Invisible Work of Accessibility: How Blind Employees Manage Accessibility in Mixed-Ability Workplaces. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility* (Lisbon, Portugal) (ASSETS '15). Association for Computing Machinery, New York, NY, USA, 163–171. doi:10.1145/2700648.2809864
- [18] Chris Callison-Burch. 2009. Fast, Cheap, and Creative: Evaluating Translation Quality Using Amazon's Mechanical Turk. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, Philipp Koehn and Rada Mihalcea (Eds.). Association for Computational Linguistics, Singapore, 286–295. <https://aclanthology.org/D09-1030/>
- [19] Joseph Chee Chang, Saleema Amershi, and Ece Kamar. 2017. Revolt: Collaborative Crowdsourcing for Labeling Machine Learning Datasets. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 2334–2346. doi:10.1145/3025453.3026044
- [20] Kyla Chasalow and Karen Levy. 2021. Representativeness in statistics, politics, and machine learning. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 77–89.
- [21] Khaoula Chehbouni, Mohammed Haddou, Jackie Chi Kit Cheung, and Golnoosh Farnadi. 2025. Neither Valid nor Reliable? Investigating the Use of LLMs as Judges. <https://arxiv.org/abs/2508.18076>

- [22] Jiahui Chen, Candace Ross, Reyhane Askari-Hemmat, Koustuv Sinha, Melissa Hall, Michal Drozdal, and Adriana Romero-Soriano. 2025. Multi-Modal Language Models as Text-to-Image Model Evaluators. <https://arxiv.org/abs/2505.00759>
- [23] Tim Connell. 2008. The Challenge of Assistive Technology and Braille Literacy. <https://www.afb.org/aw/9/1/14277> [Online; accessed 6-September-2025].
- [24] Emily Corvi, Hannah Washington, Stefanie Reed, Chad Atalla, Alexandra Chouldechova, P. Alex Dow, Jean Garcia-Gathright, Nicholas J Pangakis, Emily Sheng, Dan Vann, Matthew Vogel, and Hanna Wallach. 2025. Taxonomizing Representational Harms using Speech Act Theory. In *Findings of the Association for Computational Linguistics*. doi:10.18653/v1/2025.findings-acl.202
- [25] Amanda Coston, Anna Kawakami, Haiyi Zhu, Ken Holstein, and Hoda Heidari. 2023. A validity perspective on evaluating the justified use of data-driven decision-making algorithms. In *2023 IEEE conference on secure and trustworthy machine learning (SaTML)*. IEEE, 690–704.
- [26] Lee J Cronbach and Paul E Meehl. 1955. Construct validity in psychological tests. *Psychological bulletin* 52, 4 (1955), 281.
- [27] Samantha Dalal, Siobhan Mackenzie Hall, and Nari Johnson. 2024. Provocation: Who benefits from "inclusion" in Generative AI? <https://arxiv.org/abs/2411.09102>
- [28] Maitraye Das, Alexander J Fiannaca, Meredith Ringel Morris, Shaun K Kane, and Cynthia L Bennett. 2024. From provenance to aberrations: Image creator and screen reader user perspectives on alt text for AI-generated images. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–21.
- [29] Maitraye Das, Darren Gergle, and Anne Marie Piper. 2019. "It doesn't win you friends": Understanding Accessibility in Collaborative Writing for People with Vision Impairments. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 191 (Nov. 2019), 26 pages. doi:10.1145/3359293
- [30] Nassim Dehouche and Kullathida Dehouche. 2023. What's in a text-to-image prompt? The potential of stable diffusion in visual arts education. *Heliyon* 9, 6 (2023), e16757. doi:10.1016/j.heliyon.2023.e16757
- [31] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. <https://arxiv.org/abs/2310.00907>
- [32] Sunipa Dev, Vinodkumar Prabhakaran, Rutledge Chin Feman, Aida Davani, Remi Denton, Charu Kalia, Piyawat L Kumjorn, Madhurima Maji, Rida Qadri, Negar Rostamzadeh, Renee Shelby, Romina Stella, Hayk Stepanyan, Erin van Liemt, Aishwarya Verma, Oscar Wahltinez, Edem Wornyo, Andrew Zaldivar, and Saška Mojsilović. 2026. A Unified Framework to Quantify Cultural Intelligence of AI. <https://arxiv.org/abs/2603.01211>
- [33] Athiya Deviyani and Fernando Diaz. 2025. Contextual Metric Meta-Evaluation by Measuring Local Metric Accuracy. <https://arxiv.org/abs/2503.19828>
- [34] Lisa Egede. 2025. *Exploring Black Communities' Perceptions and Design Approaches for Building Culturally Tailored AI Systems*. Association for Computing Machinery, New York, NY, USA, 72–76. <https://doi.org/10.1145/3715668.3735629>
- [35] Maria Eriksson, Erasmo Purificato, Arman Noroozian, Joao Vinagre, Guillaume Chaslot, Emilia Gomez, and David Fernandez-Llorca. 2025. Can We Trust AI Benchmarks? An Interdisciplinary Review of Current Issues in AI Evaluation. <https://arxiv.org/abs/2502.06559>
- [36] Yannick Exner, Jochen Hartmann, Oded Netzer, and Shunyuan Zhang. 2025. AI in Disguise - How AI-Generated Ads' Visual Cues Shape Consumer Perception and Performance. doi:10.2139/ssrn.5096969
- [37] Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. 2009. Describing objects by their attributes. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. 1778–1785. doi:10.1109/CVPR.2009.5206772
- [38] Sanjana Gautam, Pranav Narayanan Venkit, and Sourojit Ghosh. 2024. From melting pots to misrepresentations: Exploring harms in Generative AI. *arXiv preprint arXiv:2403.10776* (2024).
- [39] Simret Araya Gebreegziabher, Charles Chiang, Zichu Wang, Zahra Ashktorab, Michelle Brachman, Werner Geyer, Toby Jia-Jun Li, and Diego Gómez-Zarà. 2025. MetricMate: An Interactive Tool for Generating Evaluation Criteria for LLM-as-a-Judge Workflow. In *Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '25)*. Association for Computing Machinery, New York, NY, USA, Article 22, 18 pages. doi:10.1145/3729176.3729199
- [40] Sourojit Ghosh, Pranav Narayanan Venkit, Sanjana Gautam, Shomir Wilson, and Aylin Caliskan. 2024. Do Generative AI Models Output Harm while Representing Non-Western Cultures: Evidence from A Community-Centered Approach. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* 7, 1 (Oct. 2024), 476–489. doi:10.1609/aies.v7i1.31651
- [41] Tarleton Gillespie. 2024. Generative AI and the politics of visibility. *Big Data & Society* 11, 2 (2024), 20539517241252131. doi:10.1177/20539517241252131
- [42] Luke Guerdan, Solon Barocas, Kenneth Holstein, Hanna Wallach, Zhiwei Steven Wu, and Alexandra Chouldechova. 2025. Validating LLM-as-a-Judge Systems under Rating Indeterminacy. <https://arxiv.org/abs/2503.05965>
- [43] Kanika Gupta, Monojit Choudhury, and Kalika Bali. 2012. Mining Hindi-English Transliteration Pairs from Online Hindi Lyrics. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Mehmet Uğur Doğan, Bente Maegaard, Joseph Mariani, Asuncion Moreno, Jan Odijk, and Stelios Piperidis (Eds.). European Language Resources Association (ELRA), Istanbul, Turkey, 2459–2465. <https://aclanthology.org/L12-1179/>

- [44] Rishav Hada, Varun Gumma, Adrian de Wynter, Harshita Diddee, Mohamed Ahmed, Monojit Choudhury, Kalika Bali, and Sunayana Sitaram. 2024. Are Large Language Model-based Evaluators the Solution to Scaling Up Multilingual Evaluation? *arXiv preprint arXiv:2309.07462* (2024).
- [45] Melissa Hall, Samuel J. Bell, Candace Ross, Adina Williams, Michal Drozdal, and Adriana Romero Soriano. 2024. Towards Geographic Inclusion in the Evaluation of Text-to-Image Models. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro, Brazil) (FAccT '24). Association for Computing Machinery, New York, NY, USA, 585–601. doi:10.1145/3630106.3658927
- [46] Stuart Hall (Ed.). 1997. *Representation: Cultural Representations and Signifying Practices*. Sage Publications, London.
- [47] Siobhan Mackenzie Hall, Samantha Dalal, Raesetje Sefala, Foutse Yuehghoh, Aisha Alaagib, Imane Hamzaoui, Shu Ishida, Jabez Magomere, Lauren Crais, Aya Salama, et al. 2025. The Human Labour of Data Work: Capturing Cultural Diversity through World Wide Dishes. *arXiv preprint arXiv:2502.05961* (2025).
- [48] Hamna, Gayatri Bhat, Sourabrata Mukherjee, Faisal Lalani, Evan Hadfield, Divya Siddarth, Kalika Bali, and Sunayana Sitaram. 2025. Building Benchmarks from the Ground Up: Community-Centered Evaluation of LLMs in Healthcare Chatbot Settings. <https://arxiv.org/abs/2509.24506>
- [49] Hamna, Deepthi Sudharsan, Agrima Seth, Ritvik Budhiraja, Deepika Khullar, Vyshak Jain, Kalika Bali, Aditya Vashistha, and Sameer Segal. 2025. Kahani: Culturally-Nuanced Visual Storytelling Tool for Non-Western Cultures. In *Proceedings of the 2025 ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies (COMPASS '25)*. Association for Computing Machinery, New York, NY, USA, 379–400. doi:10.1145/3715335.3735478
- [50] Emma Harvey, Emily Sheng, Su Lin Blodgett, Alexandra Chouldechova, Jean Garcia-Gathright, Alexandra Olteanu, and Hanna Wallach. 2025. Understanding and Meeting Practitioner Needs When Measuring Representational Harms Caused by LLM-Based Systems. <https://arxiv.org/abs/2506.04482>
- [51] Helia Hashemi, Jason Eisner, Corby Rosset, Benjamin Van Durme, and Chris Kedzie. 2024. LLM-Rubric: A Multidimensional, Calibrated Approach to Automated Evaluation of Natural Language Texts. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 13806–13834. doi:10.18653/v1/2024.acl-long.745
- [52] Huiguo He, Huan Yang, Zixi Tuo, Yuan Zhou, Qiuyue Wang, Yuhang Zhang, Zeyu Liu, Wenhao Huang, Hongyang Chao, and Jian Yin. 2025. Dreamstory: Open-domain story visualization by llm-guided multi-subject consistent diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2025).
- [53] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2018. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. <https://arxiv.org/abs/1706.08500>
- [54] Rachel Hong, William Agnew, Tadayoshi Kohno, and Jamie Morgenstern. 2024. Who's in and who's out? A case study of multimodal CLIP-filtering in DataComp. In *Proceedings of the 4th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–17.
- [55] Chien-Chi Hsu and Brian A. Sandford. 2007. The Delphi technique: Making sense of consensus. *Practical Assessment, Research, and Evaluation* 12, 10 (2007), 1–8. <https://openpublishing.library.umass.edu/pare/article/id/1418/> A widely cited methodological overview of the Delphi method.
- [56] Yushi Hu, Benlin Liu, Jungo Kasai, Yizhong Wang, Mari Ostendorf, Ranjay Krishna, and Noah A Smith. 2023. TIFA: Accurate and Interpretable Text-to-Image Faithfulness Evaluation with Question Answering. <https://arxiv.org/abs/2303.11897>
- [57] Mina Huh, Yi-Hao Peng, and Amy Pavel. 2023. GenAssist: Making image generation accessible. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–17.
- [58] Sadeep Jayasumana, Srikumar Ramalingam, Andreas Veit, Daniel Glasner, Ayan Chakrabarti, and Sanjiv Kumar. 2024. Rethinking FID: Towards a Better Evaluation Metric for Image Generation. <https://arxiv.org/abs/2401.09603>
- [59] Akshita Jha, Vinodkumar Prabhakaran, Remi Denton, Sarah Laszlo, Shachi Dave, Rida Qadri, Chandan K Reddy, and Sunipa Dev. 2024. Visage: A global-scale analysis of visual stereotypes in text-to-image generation. *arXiv preprint arXiv:2401.06310* (2024).
- [60] Harry H. Jiang, Lauren Brown, Jessica Cheng, Mehtab Khan, Abhishek Gupta, Deja Workman, Alex Hanna, Johnathan Flowers, and Timnit Gebru. 2023. AI Art and its Impact on Artists. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society* (Montréal, QC, Canada) (AIES '23). Association for Computing Machinery, New York, NY, USA, 363–374. doi:10.1145/3600211.3604681
- [61] Nari Johnson, Hamna Abid, Deepthi Sudharsan, Theo Holroyd, Samantha Dalal, Siobhan Mackenzie Hall, Jennifer Wortman Vaughan, Daniela Massiceti, and Cecily Morrison. 2025. Position: To Make Text-to-Image Models that Work for Marginalized Communities, We Need New Measurement Practices for the Long Tail. <https://www.microsoft.com/en-us/research/publication/position-to-make-text-to-image-models-that-work-for-marginalized-communities-we-need-new-measurement-practices-for-the-long-tail/>
- [62] Shivani Kapania, Stephanie Ballard, Alex Kessler, and Jennifer Wortman Vaughan. 2025. Examining the Expanding Role of Synthetic Data Throughout the AI Development Pipeline. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*.

- [63] Anna Kawakami, Su Lin Blodgett, Solon Barocas, Alex Chouldechova, Abigail Jacobs, Emily Sheng, Jenn Wortman Vaughan, Hanna Wallach, Amy Winecoff, Angelina Wang, Haiyi Zhu, and Ken Holstein. 2025. *Translation Tutorial: AI Measurement as a Stakeholder-Engaged Design Practice*. Retrieved January 10, 2026 from <https://drive.google.com/file/d/12qQd6ROfacYAtOQ-iihgQxslNSIKF1Pu/view>
- [64] Anna Kawakami, Jordan Taylor, Sarah Fox, Haiyi Zhu, and Kenneth Holstein. 2026. AI failure loops in devalued work: The confluence of overconfidence in AI and underconfidence in worker expertise. *Big Data & Society* 13, 1 (2026), 20539517261424164. doi:10.1177/20539517261424164
- [65] Hannah Rose Kirk, Alexander Whitefield, Paul Röttger, Andrew Bean, Katerina Margatina, Juan Ciro, Rafael Mosquera, Max Bartolo, Adina Williams, He He, et al. 2024. The PRISM Alignment Project: What Participatory, Representative and Individualised Human Feedback Reveals About the Subjective and Multicultural Alignment of Large Language Models. *arXiv preprint arXiv:2404.16019* (2024).
- [66] Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. 2023. Pick-a-Pic: An Open Dataset of User Preferences for Text-to-Image Generation. <https://arxiv.org/abs/2305.01569>
- [67] Kevin Knight and Jonathan Graehl. 1998. Machine Transliteration. *Computational Linguistics* 24, 4 (1998), 599–612. <https://aclanthology.org/J98-4003/>
- [68] Elisa Kreiss, Cynthia Bennett, Shayan Hooshmand, Eric Zelikman, Meredith Ringel Morris, and Christopher Potts. 2022. Context Matters for Image Descriptions for Accessibility: Challenges for Referenceless Evaluation Metrics. *arXiv preprint arXiv:2205.10646* (2022).
- [69] Neha Kumar, Naveena Karusala, Azra Ismail, Marisol Wong-Villacres, and Aditya Vishwanath. 2019. Engaging Feminist Solidarity for Comparative Research, Design, and Practice. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 167 (Nov. 2019), 24 pages. doi:10.1145/3359269
- [70] Anoop Kunchukuttan, Divyanshu Kakwani, Satish Golla, Gokul N. C., Avik Bhattacharyya, Mitesh M. Khapra, and Pratyush Kumar. 2020. AI4Bharat-IndicNLP Corpus: Monolingual Corpora and Word Embeddings for Indic Languages. <https://arxiv.org/abs/2005.00085>
- [71] Tony Lee, Michihiro Yasunaga, Chenlin Meng, Yifan Mai, Joon Sung Park, Agrim Gupta, Yunzhi Zhang, Deepak Narayanan, Hannah Benita Teufel, Marco Bellagente, Minguk Kang, Taesung Park, Jure Leskovec, Jun-Yan Zhu, Li Fei-Fei, Jiajun Wu, Stefano Ermon, and Percy Liang. 2023. Holistic Evaluation of Text-To-Image Models. <https://arxiv.org/abs/2311.04287>
- [72] Dawei Li, Bohan Jiang, Liangjie Huang, Alimohammad Beigi, Chengshuai Zhao, Zhen Tan, Amrita Bhattacharjee, Yuxuan Jiang, Canyu Chen, Tianhao Wu, Kai Shu, Lu Cheng, and Huan Liu. 2025. From Generation to Judgment: Opportunities and Challenges of LLM-as-a-judge. <https://arxiv.org/abs/2411.16594>
- [73] Zhiqiu Lin, Deepak Pathak, Baiqi Li, Jiayao Li, Xide Xia, Graham Neubig, Pengchuan Zhang, and Deva Ramanan. 2024. Evaluating Text-to-Visual Generation with Image-to-Text Generation. <https://arxiv.org/abs/2404.01291>
- [74] Kelly Mack, Rai Ching Ling Hsu, Andrés Monroy-Hernández, Brian A. Smith, and Fannie Liu. 2023. Towards Inclusive Avatars: Disability Representation in Avatar Platforms. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 607, 13 pages. doi:10.1145/3544548.3581481
- [75] Kelly Avery Mack, Rida Qadri, Remi Denton, Shaun K Kane, and Cynthia L Bennett. 2024. “They only care to show us the wheelchair”: disability representation in text-to-image AI models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–23.
- [76] Jabez Magomere, Shu Ishida, Tejumade Afonja, Aya Salama, Daniel Kochin, Yuehgo Foutse, Imane Hamzaoui, Raesetje Sefala, Aisha Alaagib, Samantha Dalal, et al. 2025. The World Wide recipe: A community-centred framework for fine-grained data collection and regional bias operationalisation. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. 246–282.
- [77] Daniela Massiceti, Camilla Longden, Agnieszka Slowik, Samuel Wills, Martin Grayson, and Cecily Morrison. 2024. Explaining CLIP’s performance disparities on data from blind/low vision users. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12172–12182.
- [78] J. Nathan Matias and Megan Price. 2025. How public involvement can improve the science of AI. *Proceedings of the National Academy of Sciences* 122, 48 (2025), e2421111122. doi:10.1073/pnas.2421111122
- [79] Timothy R McIntosh, Teo Susnjak, Nalin Arachchilage, Tong Liu, Dan Xu, Paul Watters, and Malka N Halgamuge. 2025. Inadequacies of Large Language Model Benchmarks in the Era of Generative Artificial Intelligence. *IEEE Transactions on Artificial Intelligence* (2025), 1–18. doi:10.1109/tai.2025.3569516
- [80] Annika Muehlbradt and Shaun K. Kane. 2022. What’s in an ALT Tag? Exploring Caption Content Priorities through Collaborative Captioning. *ACM Trans. Access. Comput.* 15, 1, Article 6 (March 2022), 32 pages. doi:10.1145/3507659
- [81] Peter Newmark. 1988. *A Textbook of Translation*. Vol. 66. Prentice Hall, New York.
- [82] Tonya Nguyen, Jean Garcia-Gathright, Hannah Washington, Alexandra Chouldechova, Hanna Wallach, and Jennifer Wortman Vaughan. 2026. Validating and Refining Generative AI Evaluations via Stakeholder Engagement. In *Proceedings of the 9th ACM Conference on Fairness, Accountability, and Transparency*.
- [83] OpenAI. 2022. DALL-E now available without waitlist. <https://openai.com/index/dall-e-now-available-without-waitlist/> Accessed: 2024-05-28.
- [84] OpenAI. 2024. DALL-E 3. <https://openai.com/index/dall-e-3/>. Accessed: 2025-09-11.

- [85] OpenAI. 2025. GPT-Image-1. <https://platform.openai.com/docs/guides/image-generation?image-generation-model=gpt-image-1>. Accessed: 2025-09-11.
- [86] OpenAI and others. 2024. GPT-4o System Card. <https://arxiv.org/abs/2410.21276>
- [87] OpenAI Developer Community Forum. 2023. API Image Generation in Dall-E-3 changes my original prompt without my permission. <https://community.openai.com/t/api-image-generation-in-dall-e-3-changes-my-original-prompt-without-my-permission/476355/2> [Online; accessed 6-September-2025].
- [88] Iroro Orife, Julia Kreutzer, Blessing Sibanda, Daniel Whitenack, Kathleen Siminyu, Laura Martinus, Jamiil Toure Ali, Jade Abbott, Vukosi Marivate, Salomon Kabongo, Musie Meressa, Espoir Murhabazi, Orevaoghene Ahia, Elan van Biljon, Arshath Ramkilowan, Adewale Akinfaderin, Alp Öktem, Wole Akin, Ghollah Kioko, Kevin Degila, Herman Kamper, Bonaventure Dossou, Chris Emezue, Kelechi Ogueji, and Abdallah Bashir. 2020. Masakhane – Machine Translation For Africa. <https://arxiv.org/abs/2003.11529>
- [89] Will Orr and Edward B. Kang. 2024. AI as a Sport: On the Competitive Epistemologies of Benchmarking. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (Rio de Janeiro, Brazil) (FAccT '24)*. Association for Computing Machinery, New York, NY, USA, 1875–1884. doi:10.1145/3630106.3659012
- [90] Mayu Otani, Riku Togashi, Yu Sawai, Ryosuke Ishigami, Yuta Nakashima, Esa Rahtu, Janne Heikkilä, and Shin'ichi Satoh. 2023. Toward Verifiable and Reproducible Human Evaluation for Text-to-Image Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14277–14286.
- [91] Lawrence Palinkas, Sarah Horwitz, Carla Green, Jennifer Wisdom, Naihua Duan, and Kimberly Hoagwood. 2013. Purposeful Sampling for Qualitative Data Collection and Analysis in Mixed Method Implementation Research. *Administration and policy in mental health* 42 (11 2013). doi:10.1007/s10488-013-0528-y
- [92] Qian Pan, Zahra Ashktorab, Michael Desmond, Martín Santillán Cooper, James Johnson, Rahul Nair, Elizabeth Daly, and Werner Geyer. 2024. Human-Centered Design Recommendations for LLM-as-a-judge. In *Proceedings of the 1st Human-Centered Large Language Modeling Workshop*, Nikita Soni, Lucie Flek, Ashish Sharma, Diyi Yang, Sara Hooker, and H. Andrew Schwartz (Eds.). ACL, TBD, 16–29. doi:10.18653/v1/2024.hucllm-1.2
- [93] Rida Qadri, Mark Diaz, Ding Wang, and Michael Madaio. 2025. The Case for "Thick Evaluations" of Cultural Representation in AI. *arXiv preprint arXiv:2503.19075* (2025).
- [94] Rida Qadri, Renee Shelby, Cynthia L Bennett, and Remi Denton. 2023. Ai's regimes of representation: A community-centered study of text-to-image models in south asia. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 506–517.
- [95] Inioluwa Deborah Raji, Emily M. Bender, Amandalynne Paullada, Emily Denton, and Alex Hanna. 2021. AI and the Everything in the Whole Wide World Benchmark. <https://arxiv.org/abs/2111.15366>
- [96] David Romero, Chenyang Lyu, Haryo Akbarianto Wibowo, Teresa Lynn, Injy Hamed, Aditya Nanda Kishore, Aishik Mandal, Alina Dragonetti, Artem Abzaliev, Atnafu Lambebo Tonja, Bontu Fufa Balcha, Chenxi Whitehouse, Christian Salamea, Dan John Velasco, David Ifeoluwa Adelani, David Le Meur, Emilio Villa-Cueva, Fajri Koto, Fauzan Farooqui, Frederico Belcavello, Ganzorig Batnasan, Gisela Vallejo, Grainne Caulfield, Guido Ivetta, Haiyue Song, Henok Biadglign Ademtew, Hernán Maina, Holy Lovenia, Israel Abebe Azime, Jan Christian Blaise Cruz, Jay Gala, Jiahui Geng, Jesus-German Ortiz-Barajas, Jinheon Baek, Jocelyn Dunstan, Laura Alonso Alemany, Kumaranage Ravindu Yasas Nagasinghe, Luciana Benotti, Luis Fernando D'Haro, Marcelo Viridiano, Marcos Estecha-Garitagoitia, Maria Camila Buitrago Cabrera, Mario Rodríguez-Cantelar, Mélanie Joutiteau, Mihail Mihaylov, Mohamed Fazli Mohamed Imam, Muhammad Farid Adilazuarda, Munkhjargal Gochoo, Munkh-Erdene Otgonbold, Naome Etori, Olivier Niyomugisha, Paula Mónica Silva, Pranjal Chitale, Raj Dabre, Rendi Chevi, Ruo Chen Zhang, Ryandito Diandaru, Samuel Cahyawijaya, Santiago Góngora, Soyeong Jeong, Sukannya Purkayastha, Tatsuki Kuribayashi, Teresa Clifford, Thanmay Jayakumar, Tiago Timponi Torrent, Toqeer Ehsan, Vladimir Araujo, Yova Kementchedjheva, Zara Burzo, Zheng Wei Lim, Zheng Xin Yong, Oana Ignat, Joan Nwatu, Rada Mihalcea, Tamar Solorio, and Alham Fikri Aji. 2024. CVQA: Culturally-diverse Multilingual Visual Question Answering Benchmark. <https://arxiv.org/abs/2406.05967>
- [97] Paul Röttger, Fabio Pernisi, Bertie Vidgen, and Dirk Hovy. 2025. SafetyPrompts: a Systematic Review of Open Datasets for Evaluating and Improving Large Language Model Safety. <https://arxiv.org/abs/2404.05399>
- [98] Olawale Salaudeen, Anka Reuel, Ahmed Ahmed, Suhana Bedi, Zachary Robertson, Sudharsan Sundar, Ben Domingue, Angelina Wang, and Sanmi Koyejo. 2025. Measurement to Meaning: A Validity-Centered Framework for AI Evaluation. <https://arxiv.org/abs/2505.10573>
- [99] Michael Saxon, Fatima Jahara, Mahsa Khoshnoodi, Yujie Lu, Aditya Sharma, and William Yang Wang. 2024. Who Evaluates the Evaluations? Objectively Scoring Text-to-Image Prompt Coherence Metrics with T2IScoreScore (TS2). In *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Eds.), Vol. 37. Curran Associates, Inc., 85630–85657. https://proceedings.neurips.cc/paper_files/paper/2024/file/9b9cfd5428153ccfbd4ba34b7e007305-Paper-Conference.pdf
- [100] Agrima Seth, Sanchit Ahuja, Kalika Bali, and Sunayana Sitaram. 2024. Dosa: A dataset of social artifacts from different indian geographical subcultures. *arXiv preprint arXiv:2403.14651* (2024).
- [101] Shreya Shankar, J.D. Zamfirescu-Pereira, Bjoern Hartmann, Aditya Parameswaran, and Ian Arawjo. 2024. Who Validates the Validators? Aligning LLM-Assisted Evaluation of LLM Outputs with Human Preferences. In *Proceedings of the 37th Annual ACM Symposium on*

- User Interface Software and Technology* (Pittsburgh, PA, USA) (UIST '24). Association for Computing Machinery, New York, NY, USA, Article 131, 14 pages. doi:10.1145/3654777.3676450
- [102] Devinder Singh and Manoj Sharma. 2009. Unfolding the Indian cultural mosaic: a cross-cultural study of four regional cultures. *International Journal of Indian Culture and Business Management - Int J Indian Cult Bus Manag* 2 (01 2009). doi:10.1504/IJICBM.2009.023547
- [103] Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2020. Participation is not a Design Fix for Machine Learning. <https://arxiv.org/abs/2007.02423>
- [104] Taylor Sorensen, Jared Moore, Jillian Fisher, Mitchell Gordon, Niloofar Mireshghallah, Christopher Michael Rytting, Andre Ye, Liwei Jiang, Ximing Lu, Nouha Dziri, Tim Althoff, and Yejin Choi. 2024. A Roadmap to Pluralistic Alignment. <https://arxiv.org/abs/2402.05070>
- [105] Stability AI. 2024. *Announcing the Open Release of Stable Diffusion 3 Medium, Our Most Sophisticated Image Generation Model to Date*. <https://stability.ai/news/stable-diffusion-3-medium> Accessed: 2025-09-11.
- [106] Stability AI. 2024. *Introducing Stable Diffusion 3.5*. <https://stability.ai/news/introducing-stable-diffusion-3-5> Accessed: 2025-09-11.
- [107] Harini Suresh, Divya Shanmugam, Tiffany Chen, Annie G Bryan, Alexander D'Amour, John Guttag, and Arvind Satyanarayan. 2023. Kaleidoscope: Semantically-grounded, context-specific ML model evaluation. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 775, 13 pages. doi:10.1145/3544548.3581482
- [108] Harini Suresh, Emily Tseng, Meg Young, Mary Gray, Emma Pierson, and Karen Levy. 2024. Participation in the Age of Foundation Models. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 1609–1621.
- [109] Annalisa Szymanski, Simret Araya Gebreegziabher, Oghenemaro Anuyah, Ronald A. Metoyer, and Toby Jia-Jun Li. 2024. Comparing Criteria Development Across Domain Experts, Lay Users, and Models in Large Language Model Evaluation. <https://arxiv.org/abs/2410.02054>
- [110] Annalisa Szymanski, Noah Ziems, Heather A. Eicher-Miller, Toby Jia-Jun Li, Meng Jiang, and Ronald A. Metoyer. 2024. Limitations of the LLM-as-a-Judge Approach for Evaluating LLM Outputs in Expert Knowledge Tasks. <https://arxiv.org/abs/2410.20266>
- [111] S. J. Tambiah. 1967. Social Change in Modern India. By M. N. Srinivas. University of California Press: Berkeley and Los Angeles, and Cambridge University Press: London, 1966. Pp. xv + 194, 40s. *Modern Asian Studies* 1, 4 (1967), 404–405. doi:10.1017/S0026749X00002687
- [112] Jordan Taylor, Ellen Simpson, Anh-Ton Tran, Jed R. Brubaker, Sarah E Fox, and Haiyi Zhu. 2024. Cruising Queer HCI on the DL: A Literature Review of LGBTQ+ People in HCI. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 507, 21 pages. doi:10.1145/3613904.3642494
- [113] The Royal National Institute of Blind People. 2025. About us. <https://www.rnib.org.uk/about-us/> [Online; accessed 6-September-2025].
- [114] Anja Thieme, Rita Faia Marques, Martin Grayson, Sidhika Balachandar, Cameron Tyler Cassidy, Madiha Zahrah Choksi, Camilla Longden, Reeda Shimaz Huda, Nicholas Ileva Kalovwe, Christina Mallon, Courtney Mansperger, Daniela Massiceti, Bhaskar Mitra, Ruth Mueni Nzioka, Ioana Tanase, Yuzhe You, and Cecily Morrison. 2026. Engaging Communities Meaningfully in Defining Disability Representation for AI Image Generation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '26)*. ACM. https://www.microsoft.com/en-us/research/wp-content/uploads/2026/01/CHI_paper_FINAL.pdf
- [115] Emily Tseng, Meg Young, Marianne Aubin Le Quéré, Aimee Rinehart, and Harini Suresh. 2025. "Ownership, Not Just Happy Talk": Co-Designing a Participatory Large Language Model for Journalism. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*. Association for Computing Machinery, New York, NY, USA, 3119–3130. doi:10.1145/3715275.3732198
- [116] Sruthi Viswanathan, Seray Ibrahim, Ravi Shankar, Reuben Binns, Max Van Kleek, and Petr Slovak. 2025. The Interaction Layer: An Exploration for Co-Designing User-LLM Interactions in Parental Wellbeing Support Systems. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 310, 25 pages. doi:10.1145/3706598.3714088
- [117] Hanna Wallach, Meera Desai, A. Feder Cooper, Angelina Wang, Chad Atalla, Solon Barocas, Su Lin Blodgett, Alexandra Chouldechova, Emily Corvi, P. Alex Dow, Jean Garcia-Gathright, Alexandra Olteanu, Nicholas Pangakis, Stefanie Reed, Emily Sheng, Dan Vann, Jennifer Wortman Vaughan, Matthew Vogel, Hannah Washington, and Abigail Z. Jacobs. 2025. Position: Evaluating generative ai systems is a social science measurement challenge. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*.
- [118] Suzanna Danuta Walters. 2003. *All the rage: The story of gay visibility in America*. University of Chicago Press.
- [119] Laura Weidinger, Maribeth Rauh, Nahema Marchal, Arianna Manzini, Lisa Anne Hendricks, Juan Mateos-Garcia, Stevie Bergman, Jackie Kay, Conor Griffin, Ben Bariach, et al. 2023. Sociotechnical Safety Evaluation of Generative AI Systems. *arXiv preprint arXiv:2310.11986* (2023).
- [120] Joseph Weizenbaum. 1976. *Computer Power and Human Reason: From Judgment to Calculation*. W. H. Freeman & Co., USA.
- [121] Etienne Wenger. 1998. *Communities of Practice: Learning, Meaning, and Identity*. Cambridge University Press.
- [122] Kimi Wenzel, Nitya Devireddy, Cam Davison, and Geoff Kaufman. 2023. Can Voice Assistants Be Microaggressors? Cross-Race Psychological Responses to Failures of Automatic Speech Recognition. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 109, 14 pages. doi:10.1145/3544548.3581357

- [123] Kimi Wenzel, Avijit Ghosh, Sachin Pendse, Stephanie Milani, Ajeet Singh, Laura Dabbish, and Geoff Kaufman. 2025. Invisible by Design? Generative AI and Mirrors of Misrepresentation. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT) CRAFT*. <https://earthling.my.canva.site/invisible> Workshop (CRAFT).
- [124] Genta Indra Winata, Frederikus Hudi, Patrick Amadeus Irawan, David Anugraha, Rifki Afina Putri, Yutong Wang, Adam Nohejl, Ubaidillah Ariq Prathama, Nedjma Ousidhoum, Afifa Amriani, et al. 2024. WorldCuisines: A Massive-Scale Benchmark for Multilingual and Multicultural Visual Question Answering on Global Cuisines. *arXiv preprint arXiv:2410.12705* (2024).
- [125] Binwei Yao, Ming Jiang, Tara Bobinac, Diyi Yang, and Junjie Hu. 2024. Benchmarking Machine Translation with Cultural Awareness. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 13078–13096. doi:10.18653/v1/2024.findings-emnlp.765
- [126] Jiayi Ye, Yanbo Wang, Yue Huang, Dongping Chen, Qihui Zhang, Nuno Moniz, Tian Gao, Werner Geyer, Chao Huang, Pin-Yu Chen, Nitesh V Chawla, and Xiangliang Zhang. 2024. Justice or Prejudice? Quantifying Biases in LLM-as-a-Judge. <https://arxiv.org/abs/2410.02736>
- [127] Meg Young. 2025. Participatory AI? Begin with the Most Affected People. *TechPolicy.Press* (February 19 2025). <https://www.techpolicy.press/participatory-ai-begin-with-the-most-affected-people/> Essay published as part of the Participatory AI Research & Practice Symposium reflections series.
- [128] Meg Young, Uphol Ehsan, Ranjit Singh, Emnet Tafesse, Michele Gilman, Christina Harrington, and Jacob Metcalf. 2024. Participation versus Scale: Tensions in the Practical Demands on Participatory AI. *First Monday* (2024).
- [129] Alice Qian Zhang, Judith Amores, Hong Shen, Mary Czerwinski, Mary L. Gray, and Jina Suh. 2025. AURA: Amplifying Understanding, Resilience, and Awareness for Responsible AI Content Work. *Proc. ACM Hum.-Comput. Interact.* 9, 2, Article CSCW033 (May 2025), 45 pages. doi:10.1145/3710931
- [130] Kexin Zhang, Elmira Deldari, Zhicong Lu, Yaxing Yao, and Yuhang Zhao. 2022. “It’s Just Part of Me:” Understanding Avatar Diversity and Self-presentation of People with Disabilities in Social Virtual Reality. In *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility (Athens, Greece) (ASSETS ’22)*. Association for Computing Machinery, New York, NY, USA, Article 4, 16 pages. doi:10.1145/3517428.3544829
- [131] Xinlu Zhang, Yujie Lu, Weizhi Wang, An Yan, Jun Yan, Lianke Qin, Heng Wang, Xifeng Yan, William Yang Wang, and Linda Ruth Petzold. 2023. GPT-4V(ision) as a Generalist Evaluator for Vision-Language Tasks. <https://arxiv.org/abs/2311.01361>
- [132] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. <https://arxiv.org/abs/2306.05685>
- [133] Naitian Zhou, David Bamman, and Isaac L Bleaman. 2025. Culture is not trivia: Sociocultural theory for cultural nlp. *arXiv preprint arXiv:2502.12057* (2025).

Appendix

This appendix contains supplementary materials for our study. We include the full evaluation rubrics for each cultural artifact (Appendix A), our complete protocols and other study materials (Appendix B), and supplemental experimental results (Appendix C).

Appendix Contents

A	Complete set of evaluation rubrics	25
B	Community workshop methodologies & study materials	27
B.1	Summary of differences between protocols	27
B.2	Extended study protocol: Assistive technologies used by the blind and low vision community	27
B.3	Extended study protocol: Indian cultural artifacts	39
C	Supplemental experimental results	43
C.1	Supplemental results: Validating the systematized concepts using annotations from workshops	43
C.2	Supplemental results: Qualitatively comparing LLM-generated vs. community-informed rubrics	45
C.3	Extended methodology: Applying rubric criteria using human and LLM annotators	48
C.4	Supplemental results: Exploring inter-annotator agreement for manual rubric application	50
C.5	Supplemental results: What does manual application of rubrics reveal about cultural representation?	52
C.6	Supplemental results: Evaluating human-MLLM agreement at applying community-informed rubrics	57

A Complete set of evaluation rubrics

As described in the main text, the systematization of cultural appropriateness for each artifact is organized around binary (Y/N) criteria. Each criteria specifies a condition that must be met in order for a depiction of the artifact to be culturally appropriate. We say that an output is culturally appropriate if all of the criteria are satisfied. While the specific criteria vary by artifact, they are organized under themes that are shared across each community.

<i>Theme</i>	<i>Criteria for guide cane</i>	<i>Criteria for braille notetaker</i>
Theme 1. The object needs to be functional as an assistive technology, and usable by someone who is blind.	<p>C1. No deformed canes.</p> <p>C2. No curved (crooked) handles.</p> <p>C3. The cane must be shaped like a long (5 foot) stick.</p> <p>C4. The body must have sections that are a white color.</p> <p>C5. There must be a tip at the bottom of the cane.</p>	<p>C1. The device must be shaped like a thin rectangular box.</p> <p>The device must have a valid braille output, to read:</p> <p>C2. The device must show braille.</p> <p>C3. All depictions of braille must be tactile (embossed). No depictions of braille on electronic screens.</p> <p>C4. Depictions of braille must be valid: arranged in cells with 3 or 4 rows, and 2 columns.</p> <p>The device must have a valid braille input, to write:</p> <p>C5. The device can have a qwerty keyboard, or a braille keyboard. A braille keyboard must have 3 or 4 keys (right), a space bar, and then 3 or 4 keys (left). These keys are positioned next to each other in a straight horizontal line.</p>
Theme 2. The object in the image should not be confused with other, more hegemonic objects, such as objects that are used predominantly by people who are sighted.	<p>C1. No wooden walking sticks.</p> <p>C2. No decorative striped patterns (candy canes).</p>	<p>C1. No depictions of notetaking as writing using a pen on paper.</p> <p>C2. No devices that are shaped like laptops with an electronic screen output.</p> <p>C3. No devices that are shaped like handheld calculators with an electronic screen output.</p> <p>C4. No devices that are shaped like manual typewriters.</p>

Table 2. **Community-informed evaluation rubrics to score AI images of a guide cane and braille notetaker.** The rubric criteria are organized under two themes. An image is defined to be culturally appropriate if *all* of the criteria are met, and culturally inappropriate otherwise.

<i>Theme</i>	<i>Chundan Vallam</i>	<i>Mridangam</i>	<i>Pallanguzhi</i>	<i>Kasavu saree</i>
Theme 1. The artifact must retain a culturally recognizable physical structure and must not resemble objects that are popular or associated with unrelated traditions or contexts.	<p>C1. It must not resemble other passenger boats (like Kerala house boats, Chinese dragon boats, long-tailed Thai boats or ferry boats).</p> <p>C2. The boat must be long and narrow.</p> <p>C3. The bow of the boat must be a plain wooden extension without decorative structures.</p> <p>C4. The stern of the boat must be a straight pointed tip angled slightly upward.</p>	<p>C1. It must not resemble other percussion instruments (like Tabla, Drum, Damaru, Dhol).</p> <p>C2. The instrument must be long, barrel-shaped, and tapered at both ends, each ending in a rounded, double-headed form, with one end slightly larger than the other.</p> <p>C3. The body of the instrument must be made out of jackwood.</p> <p>C4. There must not be intricate design or detailed patterns on the body.</p>	<p>C1. It must not resemble other board games (like Monopoly, Tic Tac toe, etc.).</p> <p>C2. The game board must be symmetrical along the length and consist of two or three rows of pits. The rows should have at least 5 pits.</p> <p>C3. The game board must be fish or rectangular in shape.</p> <p>C4. The game board must be made out of teakwood.</p> <p>C5. The pits must be circular and evenly spaced.</p>	<p>C1. It must not resemble other items like a tablecloth, Kerala Mundu or curtains.</p> <p>C2. The saree color must be off-white with a medium width (3–5 inch) woven gold border.</p> <p>C3. The saree must be made of crisp cotton fabric throughout.</p> <p>C4. The saree must not contain heavy embellishments.</p>
Theme 2. The artifact must preserve its intended functional or performative purpose as understood within its cultural context (e.g., as a game, garment, or tool).	<p>C1. Oarsmen must sit in pairs along the length of the boat. If no oarsmen are present, consider the criteria as met.</p> <p>C2. Each oarsman must use only a single paddle. If no oarsmen are present, consider the criteria as met.</p> <p>C3. The paddle must be longer and angled downward toward the water. If no oarsmen are present, consider the criteria as met.</p> <p>C4. One person must be standing at the bow or centre position of the boat. If no oarsmen are present, consider the criteria as met.</p>	<p>C1. The heads of the instrument must be stretched goat, cow or buffalo skin.</p> <p>C2. A black circular membrane must be present in the middle of both heads and must be slightly raised from the stretched skin surfaces.</p> <p>C3. The black circular membrane on the smaller end must be slightly smaller than the one on the larger end.</p> <p>C4. The instrument must have longitudinal leather straps lacing along its body connecting the two heads under high tension.</p>	<p>C1. The size of the tokens must not be too small. The tokens should be distributable by hand.</p> <p>C2. The pits must be big enough to accommodate multiple tokens.</p>	<p>C1. The saree must be shown in a way that clearly presents its pleats and drape.</p>
Theme 3. The artifact should follow culturally appropriate placement or arrangement, as practiced in traditional usage.	<p>C1. The oarsmen must be seated facing the stern. If no oarsmen are present, consider the criteria as met.</p> <p>C2. Oarsmen must wear the same attire, typically a white traditional Kerala mundu without upper garments. If no oarsmen are present, consider the criteria as met.</p>	<p>C1. The orientation and positioning of the instrument must be horizontal, lying on its length.</p>	<p>C1. The tokens can be cowrie shells or tamarind seeds.</p>	<p><i>None</i></p>

Table 3. **Community-informed evaluation rubrics to score AI images of a Chundan Vallam, Mridangam, Pallanguzhi, and Kasavu saree.** The rubric criteria are organized under three themes that are shared across the six Indian artifacts. An image is defined to be culturally appropriate if *all* of the criteria are met, and culturally inappropriate otherwise.

B Community workshop methodologies & study materials

In this section, we present extended methodological details for both the blind and low vision (Section B.2) and Indian (Section B.3) community engagements.

B.1 Summary of differences between protocols

Our three community engagements adopt different methods to engage community members in the process of systematization. We conducted our study in two phases, where we adopted a shared methodology to engage both residents of Tamil Nadu and Kerala, and a different methodology to engage BLV community members in the UK. These different methodologies reflect different cultural contexts and best practices in creating access for each community. Each engagement was conducted and facilitated by different members of our research team, who also experimented with small adaptations when implementing our shared protocol. Table 4 summarizes notable differences in how our study methodology differed across contexts. The table provides a brief justification for why we made each methodological decision.

B.2 Extended study protocol: Assistive technologies used by the blind and low vision community

B.2.1 Creating access to images. To elicit blind and low vision community members' preferences about AI images, we needed to determine how to facilitate non-visual access to these images. First, following a best practice from past work asking blind participants to evaluate AI images [57], we created alt text for each image. To encourage consistency in the amount of detail provided for each image, a blind community member on our research team created a template of important characteristics to describe (e.g., for each image of a guide cane, we always described its shape, material, and color). We include example images and their alt text descriptions below.

We additionally drew on past scholarship on *cross-ability collaborative work* in which a blind user and sighted partner work together to complete a task [9, 17, 29, 80]. While sighted strangers may misunderstand the access needs of their collaborators [17], recent studies have adopted cross-ability protocols between participants who already know each other well and have established trust and comfort working together [80]. We follow calls from Bennett et al. [9] to understand blind community members not as passive recipients of assistance, but to instead recognize their expertise in creating access throughout the collaboration. When collaborating, we emphasized the unique skills of each participant: the blind community member as the expert on how the selected artifacts worked and how they would like their community to be represented, and their sighted collaborators as capable of providing a visual perspective on what is shown in images. When responding to images, the facilitator invited participants to further discuss what is shown in each image as a pair [80].

B.2.2 Recruitment & Workshop Activities. To recruit blind and low vision individuals who currently reside in the UK, we adopted a purposive sampling approach [91]. We recruited participants from two email lists: an internal list of blind and low vision community members who had consented to receive information about future studies at our institution, and an open list for blind and low vision technology users in the UK. We asked each blind or low vision participant to invite a sighted partner of their choice to the study, following Muehlbradt and Kane [80]. Relationships between community members and their partners included friends, partners, siblings, and children. One blind community member invited a friend who is visually impaired to participate as their buddy. More information about participants is in Table 5.

Difference between methods	Methodology for blind and low vision in the UK	Methodology for South Indian states
Prompts used to generate images	We used simple prompt templates (e.g., “a photo of a guide cane”), as described in Appendix B.2.3.	Simple prompt templates that used the transliterated name for each artifact resulted in representations that were completely and totally unrelated. We used revised prompts from DALL·E 3 that included a detailed English description of each artifact, as described in Appendix B.3.3.
Composition of the generated images	Generated images that depicted artifacts in isolation (e.g., floating in an abstract liminal space). Alt text descriptions provided to participants did not include descriptions of the surrounding scene. Discussions focused on the object in isolation.	Generated images that showed artifacts in more complex and realistic scenes, such as drums resting on the ground, or boats racing in a river. Participants who responded to images often commented on the broader scene in which the artifact appeared.
Number of community members participating in each workshop	Workshops were scheduled individually with blind and low vision community members, who could invite a sighted partner, following a past practice in cross-ability research with BLV participants [80]. To understand disagreement and variance across the community, the facilitators compared findings across workshops.	Multiple community members (4–5) participated in each workshop together. Study facilitators first asked participants to respond to images individually and then discuss their decisions as a group. As a result, disagreement across participants could be surfaced and discussed in real time [11].
Number of images shown	Participants were shown either 5 or 10 images per artifact. Images were discussed one at a time, and presented by reading an alt text description. We took short breaks to prevent participant fatigue.	Participants were shown 16 total images of each artifact, sorted into 4 groups. Images were discussed one at a time.
Rating scale	Participants were asked to provide binary judgments of cultural appropriateness: 1: This image can never be shown 2: This image can be shown	Participants were asked to make decisions on a 3-point scale: 1. Can be shown 2. Needs improvement 3. Cannot be shown

Table 4. **Summary of differences between community workshop protocols.** The table summarizes key differences between the workshop methodologies used to engage the three communities. Differences in protocols reflect differing access needs, and also small changes made at the discretion of different workshop facilitators.

Each pair of participants participated in their own workshop, following Muehlbradt and Kane [80], to give community members the space to openly discuss each activity with a partner they already felt comfortable with. Workshops were conducted synchronously online between December 2024 and March 2025, were facilitated by the first author, and ranged from 45 to 90 minutes. Participants were compensated £75, and all workshop studies were approved by our institution’s ethics review board.

Workshops began by introducing the goals of the project: to help the study designers understand how community members evaluate whether an image is an appropriate representation of their culture. To ground participant

IDs	Relationship	B's Age	B's Location	Objects discussed
B1/SP1	Friends	18–34	Sheffield, England	Braille notetaker
B2/SP2	Friends	18–34	London Area, England	Guide cane, braille notetaker
B3/SP3	Parent/Child	55–74	Undisclosed	Guide cane
B4/SP4	Parent/Child	35–54	London Area, England	Guide cane, braille notetaker
B5/VIP5	Friends	35–54	Undisclosed	Guide cane
B6/SP6	Friends	55–74	Glasgow, Scotland	Guide cane, braille notetaker
B7/SP7	Siblings	55–74	London Area, England	Braille notetaker
B8/SP8	Friends	55–74	Perth, Scotland	Guide cane, braille notetaker
B9/SP9	Friends	55–74	London Area, England	Guide cane

Table 5. **Participant information about each blind and low vision community member (B) and their sighted partners (SP).** We report each community member's age and (when disclosed) location of residence at the time of study. One participant (B5) invited a visually impaired friend (VIP5) to participate as their buddy.

discussions, we introduced the study activities by providing a hypothetical scenario for how the images they were shown would be used (“A media company has collected several images of a guide cane, and they need your help to understand which of these images they should show to users”).

For each artifact, we conducted two study activities. First, we asked participant pairs to react to both the AI-generated images and real photographs selected from Step 2. Images were presented one-at-a-time. After providing the alt-text description of the image, the facilitator asked participants to share if they felt that the image could be shown to represent the artifact, or should never be shown, and why. Based on participants' responses and interests, the facilitator asked follow-up questions to prompt participants to elaborate on what exactly made a particular image a good, bad, offensive, or incorrect depiction of the artifact. When responding to images, the facilitator invited participants to further discuss what is shown in each image as a pair [80].

The last workshop activity invited participants to reflect on what they had seen so far to create a list of the most important things that need to be shown in a culturally appropriate portrayal of the object (open-ended), and provide reasons for each of their responses. This activity encouraged participants to articulate concrete visual criteria that shaped their decisions. The study facilitator asked clarifying questions that encouraged participants to reflect on whether a characteristic could vary between portrayals of the object or encourage participants to prioritize whether some characteristics were more important than others.

B.2.3 Prompt generation templates. Table 6 shows the two prompts that we used to generate images: prompting using the artifact name, and prompting using the artifact name along with a short description written by a community member on our research team. We qualitatively observed that providing a description of each artifact resulted in improved representations, but all of the generated images we reviewed still had at least one error (e.g., with the arrangement of keys on a braille notetaker or the handle shape of a cane). Images were generated using the DALL·E 3 and Stable Diffusion 3 Medium APIs.



Prompt template and example	Example image
Artifact only: <i>“A photo of a guide cane”</i>	
Artifact + artifact description: <i>“A photo of a guide cane. A guide cane, or white cane, is an assistive technology used by people who are blind. It is a collapsible lightweight cane made of aluminum.”</i>	

Table 6. **Prompt templates used to generate images to show community members.** Artifact descriptions were written by a blind community member on the research team.

B.2.4 Selected Images & Alt Text. Below, we include the final dataset of images (including AI images and real photographs) that were shown to workshop study participants with the alt text that we provided.





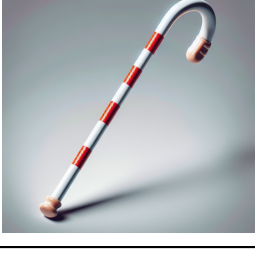
Image	Alt text description
<p>1</p> 	<p>The cane is made out of wooden material. It is a deep brown color. It has a curved grip, and straight body. The bottom part of the cane has a straight rubber tip.</p>
<p>2</p> 	<p>The cane is made out of lightweight plastic material. It is a white color, with a band of red reflective tape at the top of its body. It has a curved grip, and straight body. The bottom part of the cane has a straight rubber tip.</p>
<p>3</p> 	<p>The cane is made out of reflective metal material. It is a dark blackish brown color. The body of the cane is bent at a right angle. There is no visible handle or tip.</p>
<p>4</p> 	<p>The cane is made out of lightweight metal material. It is a light grey color, with two bands of red reflective tape. It has a straight grip, with a wrist strap. The bottom part of the cane has a round plastic marshmallow-shaped tip.</p>
<p>5</p> 	<p>The cane is made out of plastic material. It is a white color and has four wide bands of red reflective tape arranged like stripes. It has a curved grip, and straight body. The bottom part of the cane has a round mushroom shaped tip.</p>

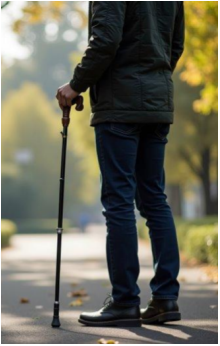




Image	Alt text description
<p>6</p> 	<p>The cane is made out of lightweight metal material. The body of the cane is black, and the handle is brown. It has a curved grip, and straight body. The bottom part of the cane has a straight rubber tip.</p>
<p>7</p> 	<p>The cane is made out of lightweight metal material. The body of the cane is a reflective red color. The cane has a crooked grip and a black handle. There is a wrist strap hanging out of the handle. The bottom part of the cane has a straight rubber tip.</p>
<p>8</p> 	<p>The cane is made out of lightweight metal material. The body of the cane is white. There are two bands of red tape on the cane. The cane has a straight body, and straight black grip at the top. There is an elastic wrist strap coming out of the grip. The bottom part of the cane has a straight rubber tip.</p>
<p>9</p> 	<p>The cane is made out of wooden material. The body of the cane is a chestnut color, and it has a black handle. It has a curved grip, and straight body. The bottom of the cane has a dark rubber tip.</p>
<p>10</p> 	<p>The cane is made of lightweight plastic material. The body of the cane is curved: there is a round handle at the top, and then two long parallel sticks coming out of each end of the handle. Each of the long sticks is a white color. At the bottom of one of the long sticks is a marshmallow-shaped red tip.</p>


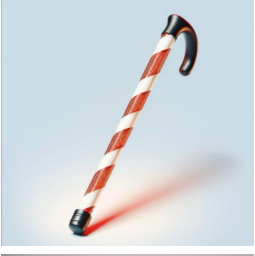
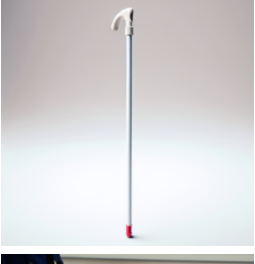
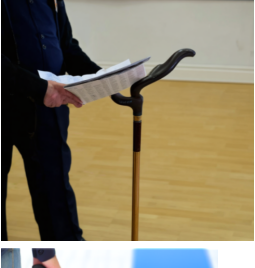
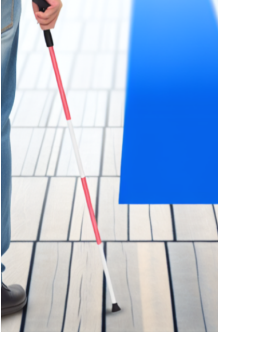
Image	Alt text description
<p>11</p> 	<p>The cane is made of lightweight metal material. The body of the cane is white. It has a straight body, and no visible handle or grip. One end of the cane has a black wrist strap. The other end of the cane has a round marshmallow-shaped tip. The body of the cane is divided into three sections by three grey joints.</p>
<p>12</p> 	<p>The cane is made of lightweight plastic material. The cane's body has wide bands of red reflective tape, arranged like stripes. The cane has a straight body and curved black handle. The bottom of the cane has a straight black tip.</p>
<p>13</p> 	<p>The cane is made of lightweight plastic material. The body of the cane is white. It has a straight body, and a handle that curves downwards. The bottom of the cane has a straight red tip.</p>
<p>14</p> 	<p>The cane is made of lightweight plastic material. The body of the cane is a gold color. The cane has a straight body, and a black curved handle. The handle has two distinct pieces that stick up at different angles at the top of the cane. The handle is irregularly curved.</p>
<p>15</p> 	<p>The cane is made of lightweight metal material. The body of the cane is white, with two wide bands of red reflective tape. The top of the cane has a straight black grip.</p>


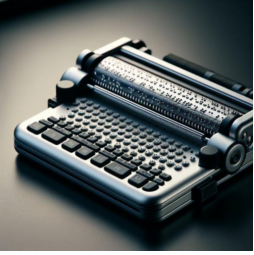



Image	Alt text description
<p>1</p> 	<p>There is a thin rectangular electronic device. The top surface of the device has a display on top, and rows of buttons below it. The display is a small electronic screen. The buttons have tactile markings on them that resemble Braille. The top surface also has a circular speaker. The sides of the device have additional ports and buttons.</p>
<p>2</p> 	<p>The device is a rectangular shape, with a roller cylinder on top of its surface where paper can be inserted. The roller has tactile markings on its surface that resemble Braille. Below the roller on the top surface, there are four rows of circular keys. Each row has about 15 keys. Below that row, there are 3 rows of larger rectangular keys.</p>
<p>3</p> 	<p>There is a device shaped like a folding laptop computer, with an electronic screen display on top and a keyboard on the bottom. The screen of the device is displaying rows of dots that resemble Braille. There are five "lines" of Braille stacked on top of each other. The keyboard of the device resembles a qwerty keyboard. There are five rows of keys, buttons at the top, and a space bar at the bottom.</p>
<p>4</p> 	<p>There is a rectangular electronic device. The top surface of the device has a display on top, and rows of buttons below it. The display has tactile markings that resemble Braille. It is showing two "lines" of Braille. Below the display, there are two stacked rows of five circular buttons. The device has other buttons on its surface, for example, that resemble a volume control. The sides of the device have additional ports and buttons. For example, one side appears to have 7 circular input ports.</p>
<p>5</p> 	<p>The image shows a paper notebook and a hand holding a pen. The notebook is open to a page. Two lines of ink dots are written on the page.</p>









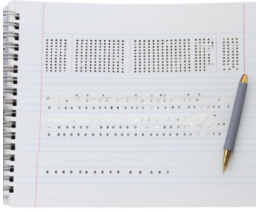

Image	Alt text description
<p>6</p> 	<p>The image shows a rectangular device. The top surface of the device has a dark display. The display has many small white tactile dots, arranged in 20 rows and 40 columns. There are no additional buttons on the sides of the device.</p>
<p>7</p> 	<p>There is a device shaped like a folding laptop computer, with an electronic screen display on top and a keyboard on the bottom. The screen of the device is displaying rows of dots that resemble Braille. There are four distinct "lines" of dots, where each "line" has about 6 rows. The keyboard of the device resembles a qwerty keyboard. There are five rows of keys, buttons at the top, and a space bar at the bottom.</p>
<p>8</p> 	<p>There is a thin rectangular electronic device. The top surface of the device has buttons on top, and a tactile display at the bottom. At the top of the device, there are eight round keys arranged in a curved pattern. There are four keys on the left, four keys on the right, and a space bar in between them. There are three additional buttons on each side of the space bar. The display has 20 cells, each of which has 4 rows, and 2 columns of tactile dots. On each side of the display, there are two buttons. The side of the device has two ports and two more small buttons.</p>
<p>9</p> 	<p>There is a thin rectangular electronic device. The top surface of the device has several buttons on top, and a tactile display at the bottom. The top of the device has one row of about 15 small circular buttons. Below that, there are two more rows of 8 circular buttons, stacked on top of each other. The display has 3 rows, and around 25 columns, of small metal circular pins that are sticking out of the device.</p>
<p>10</p> 	<p>There is a thin rectangular electronic device. It is quite wide, and not very long. The top surface of the device has a tactile display. There is no keyboard. The display has 32 cells, where each cell has 4 rows and 2 columns of small tactile dots. Each cell has a small black button above it. Next to the display, there are three circular buttons.</p>

Image	Alt text description
<p>11</p> 	<p>There is a rectangular electronic device. The top surface is shaped like a handheld calculator. The top surface of the device has a tactile display on top, and rows of buttons below it. The display has small metal tactile dots. They are irregularly arranged in two rows, and around 15 columns. Below the display, there are four rows and five columns of small oval-shaped buttons.</p>
<p>12</p> 	<p>The device is a rectangular shape, with a roller cylinder on top of its surface. Immediately above and below the roller, the surface of the device has several tactile markings that resemble braille. Below that on the top surface, there are rows of circular keys that resemble a qwerty keyboard. There are two rows of keys marked with letters, and below that there is a space bar and additional keys.</p>
<p>13</p> 	<p>There is a wide, thin electronic device. The top surface of the device has a braille display, and several buttons. At the top of the device, there is a tactile display. The display has 20 cells, each of which has 4 rows, and 2 columns of tactile dots. There are three round buttons to the left and right of the display. Below that, there is one row of eight round keys arranged in a curved pattern. There are four keys on the left, four keys on the right, and a space bar below them.</p>
<p>14</p> 	<p>The image shows a paper notebook with an ink pen resting on it. The notebook is open to a page. Several lines of ink dots have been written on the page. The top of the page has around 100 small dots, arranged in around 10 rows and 40 columns. The bottom of the page has four rows of larger ink dots. The page looks a bit crinkled, like several dots have been removed.</p>
<p>15</p> 	<p>The image shows a thin electronic device. The top surface of the device has an electronic screen, displaying the Google homepage. The bottom of the device has a tactile display. The display has 28 cells, each of which has 4 rows, and 2 columns of tactile dots. There is a button at either end of the display.</p>

B.2.5 Workshop Study Protocol. Workshops with blind and low vision community members began by the facilitator introducing themselves, and inviting the participant pair to introduce themselves. We provide the facilitator's script below:

Today, we want to understand how you would like different assistive technologies to appear in AI images.

Our goal as researchers is to learn from both of your expertise and past experiences with two assistive technologies as someone who uses them or as someone who has observed their usage.

The [first/second] object we're going to discuss, is a [OBJECT].

- If only one participant is familiar with the object: I saw from the survey that [NAME] is less familiar with [OBJECT]? Could you share more what you mean by that?

For the activity, we're going to discuss what you both think about some images. This activity is going to be based off a scenario, that I'll introduce now:

SCENARIO: "A media company has created a service to provide images for users who are making slide decks. The service works like this: every time the user asks for an image, the company generates 10 images, and then shows the user 4. The media company has collected several images of a [OBJECT]. The media company needs your help to understand which of these images they should show to users, and which images should never be shown to users."

I'll pause: Any questions about the scenario?

The company has gathered 10 total images of a [OBJECT] for us to discuss together today.

Reacting to images. For the activity today, I am going to screen share a slide deck that has the 10 images. I'm going to go through the images one-by-one and ask you both to answer some questions for each image. I'll begin by providing a basic description (some alt text) for each image. If anything is unclear about the images from the alt text I've provided, you can also ask for clarification from your partner.

You may notice that the [OBJECT] in each image is in a different scene: for example, some images show people using a [OBJECT], versus others just depict the [OBJECT] on a plain background. When answering the questions, we'd like for you to focus only on the [OBJECT] – not the scene.

Questions to ask for each image:

- Would you tell the company that this image can be shown, or should be never shown to users? Follow-up probes:
 - Why would you (not) show it?
 - What about this image is good/bad?
 - Is there something about this image that makes it an offensive or harmful depiction?
- Is this image a correct depiction of a [OBJECT]? Follow-up probes:
 - What about it makes it (in)correct?
 - Why do you think this one is OK to show, even if it is incorrect?

Summarizing all of the images: OK, so far we've selected X images that we think are not OK to show to users: [read]

- Now that we've discussed them all, are there any that we want to add to this final list?
- Are there any images that you think actually might be OK to show, and why?

Identifying important characteristics & visual criteria. Given everything we've seen, what do you think are the most important things that need to be shown for the [OBJECT] to be depicted correctly? Probes:

- Do any of the things we talk about look different for different types of [OBJECT]? Is it possible for the [thing] to vary? What could you change about [this object] for it to still be correct?
- What about [characteristic]? Would it be on your list at all?

- So far, we've listed [these X things] that are important to representing a [OBJECT]. Are there things on this list that are more important than others?

Beyond the 10 images we've looked at today, the company is interested in understanding some general principles about what makes an image good or bad to show to a user. What advice would you give to the company? What is it that makes an image of a [OBJECT] inappropriate to show?

B.3 Extended study protocol: Indian cultural artifacts

B.3.1 Recruitment & Workshop Activities. We adopted a focus group methodology where we invited participants from the same state to participate in synchronous workshops together. Past research has demonstrated the value of facilitating deliberative group discussions to surface (and clarify) points of tension and disagreement among participants [75, 94, 94]. We conducted one focus group per each artifact, where each focus group was facilitated by a community member from the relevant state. Workshops were conducted virtually and lasted 120 minutes. Participants were compensated with an Amazon voucher worth 500 INR.

State	Artifact	Number of participants	Gender ratio (M:F)	Age Range
Tamil Nadu	Pallanguzhi	4	2:2	18–34
Tamil Nadu	Mridangam	5	1:4	18–44
Kerala	Kasavu saree	4	1:3	18–54
Kerala	Chundan Vallam	4	3:1	25–44
Total		17	7:10	18–54

Table 7. **Participants in workshops organized around each Indian cultural artifact.** Workshops were conducted synchronously, and a single artifact was discussed in each workshop. Eligible participants only participated in a single workshop, and were assigned to artifacts based on their availability and expertise.

To determine eligibility, we decided to scope our study to individuals who have spent at least several years residing in Kerala or Tamil Nadu, and would therefore be likely to be familiar with cultural artifacts of the regions. To recruit participants for workshops, we aimed to capture a diverse sample of individuals who identified as being from each state. Our research team disseminated a call to participate in research by disseminating a screener survey circulated on the research team’s personal social networks, such as X.com and WhatsApp. Inclusion criteria were that participants must speak English, reside in India, and have spent at least several years residing in Kerala or Tamil Nadu. Participants also self-reported their familiarity with each cultural artifact. We identified 23 eligible participants from our initial screener survey, and invited 9 participants per state to participate. The final sample of participants who ultimately participated, along with their ages and gender identities, is summarized in Table 7.

Like the BLV workshops, the focus group activities invited community members to react to images to map the broad background concept of cultural appropriateness to specific visual characteristics. Participants were introduced to the goals of the project using a scenario where a company is exploring the use of AI-generated images in a tourism advertisement. Participants were invited to join an online board on FigJam, a collaborative web application where multiple users can take notes. In the first activity, participants were presented with 16 total AI-generated images of the artifact (4 per each of the groups from Step 2) displayed on the whiteboard. For each AI-generated image, participants were asked to rate whether they felt the image (1) cannot be shown, (2) needs improvement, or (3) can be shown as a representation of the artifact. This three-point scale provided participants with more flexibility to note when they were uncertain about a particular image. We first invited participants to evaluate each image individually by sharing their ratings in the meeting room chat. Participants were then invited to share and discuss their reactions to any of the images and justifications for their rating decisions with the group.

After the first activity, the facilitator led a group discussion inviting participants to rank the groups of images, from most to least preferred. Participants whose ratings or rankings disagreed were encouraged to understand each other’s perspectives to see if they could reach a consensus, but reaching consensus was not required.

The last study activity invited participants to reflect on if there were any varieties of the artifact that they were familiar with that had not yet been discussed. Participants were encouraged to search the Web for photographs



Fig. 6. **Stable Diffusion 3 generates depictions of unrelated cultural artifacts and scenes when given simple transliterated prompts. Depictions improve when images are generated using DALL-E 3 revised prompts instead.** The images on the left were generated with the simple transliterated prompt “A photo of a Chundan Vallam”. Instead of producing depictions of a boat, the generated images show unrelated depictions of foods and buildings, indicating that the models may struggle to interpret transliterated artifact names. The images on the right, which picture boats floating down a river, were generated by Stable Diffusion 3 when provided with a more descriptive revised prompt.

that they felt were or were not representative of the artifact. Participants discussed whether they felt each image was an appropriate cultural representation as a group.

B.3.2 Generating images. The names of each Indian cultural artifact come from the local language of each state (Tamil or Malayalam), which are written using a non-Roman script, and do not have a direct English translation. Existing text-to-image models have little to no support for these local scripts, so we followed past research and used the *transliterated* name for each object, a sound-preserving transcription of the word in Roman script [43, 67, 100]. When we used basic prompt templates with these transliterations (e.g., “a photo of a Chundan Vallam”) to generate images, we found that many of the generated images bore very little relevance to the artifacts at hand. For example, when we prompted Stable Diffusion 3 to create images of a Chundan Vallam (racing boat), it instead created images of South Asian homes, people, food, and scenes. However, images created with DALL-E 3 were more likely to create depictions that captured the essence of each artifact, due to DALL-E’s automated prompt revision feature, which rewrote our basic input prompts to include detailed English descriptions of each artifact [87]. To increase the quality of the Stable Diffusion images, we input these DALL-E revised prompts instead of using basic prompt templates. We include example DALL-E revised prompts for each artifact in the appendix.

As in the BLV context, community members on the research team sorted an initial set of 60 generated images (30 per model) into 4 groups that shared common visual characteristics so that we could prioritize getting participants’ feedback on meaningfully different depictions. The final set of AI-generated images differed from those shown to BLV participants in that they were more likely to picture each artifact within a detailed scene (e.g., of rowers paddling a racing boat down a river), in contrast to the images shown to BLV participants, which showed each artifact juxtaposed on an empty or abstract background.

B.3.3 DALL-E revised prompts. When we used the basic “artifact only” prompt template from Table 6 to generate images of each Indian artifact, we found that Stable Diffusion 3 created images that were completely unrelated to the artifact, instead of displaying representations of unrelated objects, e.g., of foods or buildings (Figure 6).

Rather than elicit community members’ feedback on images that so clearly contain errors of representation, we decided to revise our prompting strategy. To do this, we stored prompts created using DALL-E’s automated prompt revision feature, which rewrote our basic input prompts to include detailed English descriptions of each artifact [87]. To increase the quality of the Stable Diffusion images, we input these DALL-E revised prompts

instead of using basic prompt templates. We include example DALL·E 3 revised prompts for each of the Indian cultural artifacts below.

Original prompt: “a photo of a [artifact]”

- Pallanguzhi: “A realistic photograph of a traditional Pallanguzhi board on a rustic wooden table. The wooden board features intricate carvings, with two rows of seven circular pits, filled with vibrant, colorful seeds or shells used as counters. The scene is illuminated by warm natural lighting, which highlights the detailed craftsmanship of the board and the vivid hues of the seeds. Surrounding elements include traditional Indian decor, such as a brass lamp emitting a soft golden glow and a vibrant, patterned textile cloth draped in the background. The overall atmosphere of the image exudes nostalgia and serenity, celebrating the cultural heritage of this ancient Indian board game.”
- Mridangam: “A high-resolution image of an Indian mridangam, a traditional percussion instrument, placed on an ornately decorated cloth featuring intricate patterns in rich shades of gold and red. The wooden body of the mridangam is shown clearly, with its leather straps and detailed craftsmanship prominently visible. The setting is a warm, softly-lit indoor space with ambient light creating gentle shadows. The lighting emphasizes the mridangam as the centerpiece of the composition, exuding cultural beauty in a serene and elegant atmosphere.”
- Chundan Vallam: “A realistic photograph of a traditional Chundan Vallam, also known as a snake boat, floating on the still waters of Kerala’s serene backwaters. The handcrafted, dark wooden boat is long with a slender shape and an elegantly rising, pointed bow. Intricate carvings glimmer on the boat, reflecting the bright sunlight. In the backdrop, there’s a dense wall of lush, green coconut trees and tropical vegetation along the riverbanks. The sky is a clear, vibrant blue, contrasting beautifully with the greenery. The calm water creates a perfect mirror image of the boat and the surrounding landscape, enhancing the tranquil atmosphere of the scene.”
- Kasavu saree: “A serene and elegant presentation of an Indian Kasavu saree neatly folded on a rustic wooden table. The saree is predominantly white with a golden border and features intricate golden zari work embellishing its border and pallu. The scene is illuminated with warm, golden lighting, emphasizing the cultural heritage tied to Kerala. The background is softly blurred and plain to maintain a minimalist aesthetic. Near the saree, a small vase of fresh, simple flowers adds a delicate touch. The natural texture of the wooden table reinforces the timeless charm and grace of the scene, creating an artistic, warm-toned environment.”

B.3.4 Focus Group Protocol. Focus groups began by the facilitator introducing themselves, and inviting the group of workshop participants to introduce themselves. We provide the facilitator’s script below:

We’re working on a research project that involves evaluating AI-generated images of cultural artifacts of different communities – like traditional objects, clothing, foods, etc. While AI technologies are advancing rapidly, we have discovered that when generating images of commonly known cultural artifacts from India, examples often struggle to represent their cultural significance. As a cultural expert from your community, your perspective is very important to us. Our goal is to learn from your expertise and lived experiences with the cultural artifacts as someone who is from that cultural community. We have shortlisted a cultural artifact [OBJECT] that is well known to the community. The aim of this discussion is to understand how community members define the important visual characteristics of an artifact – those that must be included to make visual representations acceptable and respectful across the community. This informs our evaluation of T2I (text-to-image) generated visuals. The insights are used to improve generated visuals to ensure they reflect culturally grounded and meaningfully nuanced depictions. Some of the AI-generated images may not represent the artifact accurately or respectfully. If an image feels inappropriate or upsetting, we sincerely apologize. Please feel free to point it out – this feedback is important to help us improve.

Reacting to images. On our screen, you should see sixteen AI images of the [OBJECT]. Our team has sorted these images into four different groups that have something in common.

In this part of the study, we'll do an activity where we want you to rate each image. We'll discuss each image one-at-a-time, and are interested in knowing what everyone thinks individually before discussing each image as a group. You can message your thoughts in the chat or unmute and speak them.

There are no right or wrong answers — your unique insights matter most!

- Do you think that this image can be shown, needs improvement, or cannot be shown?
- What exactly is good or bad about the image that influences your rating?
- I noticed that we aren't sure or we disagree about [image]. Can each of you discuss why you gave your rating?

Ranking groups of images. Now that we've seen all four of these groups of images, our next goal is to rank these four groups, from best to worst.

- Which of these four groups do you think does the best job of representing the [OBJECT], and which is worst?
- Why is that group better or worse than the other one?

Providing reference images. Are there any variants or varieties of the artifacts that the model tried capturing, or didn't have knowledge of? You can feel free to browse the Internet on your devices to share photos of different versions of the [OBJECT] that we haven't discussed yet.

- Do you think that this version of the [OBJECT] would be okay to show? Why or why not?

C Supplemental experimental results

This appendix presents extended descriptions of our experimental design methodology and includes supplemental experimental results.

- We begin by presenting supplemental results that analyze the content of our community-informed rubrics (RQ1), such as a lightweight validation of our final rubric criteria (Section C.1) and a qualitative analysis that compares community-informed rubrics to those generated by an LLM (Section C.2).
- We next provide an extended description of how we *operationalized* our rubrics (RQ2) to score a new dataset of images using both human and MLLM annotators (Section C.3).
- We then present extended analyses of inter-annotator agreement among human rubric annotators (Section C.4), qualitative trends from manually scoring T2I model outputs (Section C.5), and additional results comparing human and MLLM rubric application (Section C.6).

C.1 Supplemental results: Validating the systematized concepts using annotations from workshops

Methods. Our research team manually (by hand) applied our evaluation rubric to get final labels of cultural appropriateness for each image that was shown in workshops. The research team (first and second authors) manually annotated whether each systematization criterion was met for each image. Each rubric was applied by the research team member who facilitated workshop engagements and led the creation of the rubric, to ensure its consistent application across the dataset. Each systematization criterion corresponds directly to a concrete visual indicator in an image (e.g., “does the guide cane have sections that are white in color?”). If an image does not contain enough detail to assess whether a criterion is met (e.g., if an image is cropped so that a particular feature is not visible), we default to assign a positive label (e.g., that the criterion is met) as such an image may be a potentially valid representation.

We analyze participants’ ratings to come up with a single “majority label” of cultural appropriateness for each image. In workshops, each image was shown to multiple participants, who provided binary annotations of whether they felt the image was an appropriate or inappropriate depiction of each artifact. For Indian participants, who used a three-point rating scale, we binarize ratings by treating images scored “1 – Cannot show” as inappropriate; and images scored “2 – Needs improvement” or “3 – Can show” as appropriate. We aggregate participant labels using a majority vote. In cases of ties, we label contested images as inappropriate.

	Number of images	Agreement	% contested	Community base rate	Measure base rate
All BLV artifacts	20	0.90	0.35		
Guide cane	10	0.90	0.30	0.20	0.30
Braille notetaker	10	0.90	0.40	0.30	0.20
All Indian artifacts	64	0.83	0.32		
Pallanguzhi	16	0.81	0.63	0.31	0.25
Mridangam	16	0.88	0.25	0.12	0.00
Chundan Vallam	16	0.81	0.13	0.25	0.06
Kasavu saree	16	0.81	0.25	0.50	0.31

Table 8. **Comparing community annotations from workshops to scores obtained from applying evaluation rubrics.** We find that community annotations agree with the labels of cultural appropriateness assigned by our rubrics for over 80% of the images shown to workshop participants. We report the “base rate” (the proportion of images with positive labels) for both the majority labels assigned by community members and labels assigned by our measures.

Results. Table 8 compares rubric-based annotations applied by the research team with cultural appropriateness ratings from workshop participants. We report agreement, defined as the proportion of images for which the rubric's final binary label matches the participants' majority label. To contextualize these agreement rates, we also report the proportion of images where participants' labels of cultural appropriateness disagreed with each other (% contested), as well as the overall proportion of images labeled as culturally appropriate by both.

We find that our rubric-based labels agree with participants' majority judgments for over 80% of the images shown in workshops, suggesting close alignment with community members' evaluations on this dataset. At the same time, participants' judgments themselves exhibit considerable disagreement, ranging from 13% to 63% of images across artifacts, underscoring the inherently subjective and contested nature of cultural representation. This level of intra-community disagreement suggests that perfect alignment with participant ratings is neither achievable nor necessarily desirable, as any formalized measure will inevitably reflect particular interpretations within a heterogeneous community rather than a single, unified standard.

While this level of agreement is in part expected given that the rubric was derived from participant feedback, the result nonetheless serves as an important validation step. In particular, it demonstrates that qualitative, free-text feedback from community discussions can be systematically translated into concrete visual criteria and recomposed into an evaluation rubric that reproduces participants' judgments at scale.







Correcting errors Braille notetaker		Clarifying criteria Braille notetaker		Adding criteria Mridangam	
LLM 	Includes a screen or output display for visual feedback.	LLM 	Displays rows of raised dots for braille reading.	LLM 	...
Ours 	No devices with an electronic screen output.	Ours 	The braille must be arranged in valid cells: dots in 3 or 4 rows, 2 columns.	Ours 	A black circular membrane must be present on both drumheads.

Fig. 7. **Community-elicited rubrics differ meaningfully from those generated by off-the-shelf LLMs.** Our rubrics differ from LLM-generated rubrics in three ways, each illustrated using an example (Appendix C.2). First, LLM-generated rubrics can include factual or interpretive errors that reflect misunderstandings of the artifact (e.g., whether a braille notetaker should have a screen). Second, our rubrics provide culturally grounded clarifications for features that LLM-generated rubrics leave underspecified. Third, our rubrics include additional criteria that LLM-generated rubrics omit, such as the black circular membrane on a Mridangam drum.

C.2 Supplemental results: Qualitatively comparing LLM-generated vs. community-informed rubrics

Using an LLM to generate a rubric. We adapt the prompt that Szymanski et al. [109] provide an LLM to generate rubrics, making small revisions to instruct the model to identify visual criteria that can be used to evaluate images. We provide each prompt to GPT-4o and sample a single rubric per artifact. The final set of generated rubrics contain 5-6 criteria that describe visual characteristics that must be present in a culturally appropriate depiction of each artifact.

System Prompt. You are a helpful and precise assistant that can create binary evaluation criteria to evaluate images of cultural artifacts. Your task is to generate evaluation criteria for assessing whether an image contains a culturally appropriate depiction of an object. Each criterion should be a statement in which you would answer true/false. The criteria should describe the most important visual characteristics that should be present or absent in a correct depiction of the object. The criteria should not be in the form of a one sentence statement, not a question. You should return your final answer as a valid JSON object.

User Prompt. Create evaluation criteria for the given prompt instruction:
"A photo of a guide cane"

Analysis methodology. We follow Szymanski et al. [109] to conduct a simple qualitative comparison of LLM-generated and human-generated (community-informed) rubrics. Because our dataset only contained six rubrics with a small set of criteria, our research team was able to manually review and assess the complete set of LLM-generated criteria. To conduct the comparison, the first author annotated each LLM rubric in collaboration with community members on the research team, to highlight key differences between what was articulated in each LLM rubric, versus our own understandings of each artifact, as shaped by what we learned from workshops and also our own lived understandings of what these artifacts should be. For each criterion, we assessed if it was related to or overlapped with another criteria in our rubrics. We paid attention to places where each LLM rubric diverged from views that were articulated by the community, or our own understandings of each artifact. The final result of this analysis included an assessment of each individual criterion, and also broader trends in what was, or was not, captured by these LLM-generated rubrics, summarized in Figure 7.

C.2.1 Results. We share our research team’s annotations of the rubric criteria and provide qualitative descriptions of how each rubric differs from our community-elicited rubrics for each artifact.

- C1: The image contains a long, slender stick resembling a cane.
- C2: The cane is primarily white with a distinctive red band near the bottom.
- C3: The handle of the cane is designed for comfortable grip, often featuring ergonomics or a looped strap.
- C4: The tip of the cane is designed for ground contact, typically cylindrical or spherical.
- C5: There are no extraneous decorations or embellishments that deviate from its functional design.

Fig. 8. **Annotated LLM-generated rubric for a guide cane.** While generally providing an accurate description of a guide cane, the rubric misses several key details. The rubric does not provide a complete description of the straight handle shape of a cane (C2), a feature that is of critical importance to the community. In workshops, we learned that a band of red tape on a cane’s body is often a visual signifier that the user is deaf-blind (C3), and thus should not be required for a culturally appropriate depiction.

- C1: The image contains a device resembling an electronic notepad with a braille display.
- C2: The image displays rows of raised dots for braille reading.
- C3: The device in the image includes a qwerty keyboard or braille input keys.
- C4: The image shows a size appropriate for handheld or portable use.
- C5: The device includes a screen or output display for visual feedback.

Fig. 9. **Annotated LLM-generated rubric for a braille notetaker.** The rubric criteria include both inaccurate descriptions of a braille notetaking device, and do not include descriptive details about valid depictions of braille. A braille notetaker does not resemble a notepad (e.g., it does not include a writing device such as a pen), and instead resembles a slim rectangular box (C1). The rubric does not provide a description of braille beyond “raised dots”, which is underspecified as braille must be arranged in valid cells to be readable (C2). Similarly, the rubric does not describe the unique layout of “braille input keys” (C3). Many braille notetaking devices do not have a screen or visual display, as such displays are inaccessible to blind and low vision community members (C5).

- C1: The image contains a wooden or traditional material board with two rows of pits.
- C2: The board has seven pits in each row, totaling fourteen pits.
- C3: The pits are uniform in size and symmetrically arranged on the board.
- C4: The image includes small shells, seeds, or beads used as playing pieces, visible either in the pits or nearby.
- C5: There is a clear depiction of traditional or cultural motifs or designs on the board or surrounding area.

Fig. 10. **Annotated LLM-generated rubric for Pallanguzhi.** The rubric generally provides an accurate description of the most important characteristics of the a Pallanguzhi board, with two differences from the community-elicited rubric. (C1) Community members clarified that the color of the wood is important, and that Pallanguzhi boards are traditionally made of a deep-brown teakwood. (C2) The number of pits in each row of a Pallanguzhi board can vary between 5 and 7. They also emphasized the size of the playing pieces or tokens, noting that they should not be too small and should be similar to tamarind seeds or cowrie shells (C4).

- C1: The image includes an elongated drum with a barrel-shaped body.
- C2: The drum has **two drumheads**, one on each side, with visible attaching straps or cords.
- C3: There is an absence of any modern or electrified components attached to the drum.
- C4: The **wood** of the drum body shows a polished, traditional finish, **possibly with minimal decorative carvings**.
- C5: The image portrays traditional materials such as leather or animal skin for the **drumheads**.
- C6: There are no inappropriate or non-cultural symbols or markings on the drum.

Fig. 11. **Annotated LLM-generated rubric for a Mridangam.** The rubric lacks many of the critical details that distinguish the Mridangam from related drums and percussion instruments. One significant omission is the black circular membrane that must be present on both drumheads, a key feature that contributes to the timbre of the drum (C5). One drumhead is often slightly larger than the other (C2). The Mridangam should not be depicted with decorative patterns (C4). The rubric lacks details about the characteristic horizontal orientation of the drum, which must be played on its length.

- C1: The saree in the image has a white or cream base color.
- C2: The saree features a distinct golden border.
- C3: The image shows the saree being worn in a traditional drape.
- C4: The texture of the fabric appears to be cotton **or silk**.
- C5: The overall presentation of the saree respects cultural context and dignity.

Fig. 12. **Annotated LLM-generated rubric for a Kasavu saree.** The rubric generally provides an accurate description of a Kasavu saree, demonstrating substantial overlap with the community-elicited rubric. However, community members were clear that the material must be cotton and not silk (C4).

- C1: The depiction shows a long and narrow wooden boat traditionally used in Kerala, India.
- C2: The image includes details such as **oars, indicating its use as a rowed vessel**.
- C3: The boat is adorned with distinct decorative elements characteristic of Chundan Vallams, such as colorful flags or cloth.
- C4: The scene includes a setting typical of Kerala, such as backwaters or rivers.
- C5: **The image presents the boat with a curved, pointed prow, a common feature of this type of vessel.**
- C6: The boat is **shown with a crew of rowers**, if visible, in traditional attire.
- C7: The photograph maintains cultural context by not altering or modernizing the appearance of the boat inappropriately.

Fig. 13. **Annotated LLM-generated rubric for Chundan Vallam.** The rubric criteria cover the general structure of the Chundan Vallam but do not specify its defining features. In particular, they omit details about the oar structure and handling (C2); community members specified that the oars should be long, angled downward toward the water, and that each oarsman must use a single oar. The rubrics also do not specify the distinct characteristic of the stern being a straight, pointed tip (C5). Additionally, they lack guidance on the seating position of rowers; as emphasized by community members, when rowers are visible in images, they should be seated in pairs and face towards the stern (C6).

C.3 Extended methodology: Applying rubric criteria using human and LLM annotators

In this section, we provide an extended description of the methods that we adopted to operationalize (apply) our evaluation rubrics to score a new dataset of images in Section 4.2. We first describe how we generated a dataset of 50 images per each artifact. We then provide an extended description of the methods that we used to apply the rubrics manually (by collecting annotations from a member of our research team), and using a multimodal LLM-as-a-judge pipeline.

C.3.1 Generating images. We generated 10 images each using five image generation models that achieved state-of-the-art performance at the time of our study. These include the two models from our case studies [83, 84, 105] and three newer models, including OpenAI’s latest, GPT-Image-1 [14, 85, 106]. When generating images with GPT-Image-1 or DALL·E 3, we used the simple prompt “A photo of a {artifact}”³. For the Stable Diffusion and Flux models, we expanded this prompt with more detailed descriptions of each artifact to improve the depictions’ quality as these models performed poorly out-of-the-box. For the assistive technologies, we used prompts containing a description written by a community member on our research team (Figure 6). For the Indian cultural artifacts, we used revised DALL·E 3 prompts as detailed in Appendix B.3.3.

C.3.2 Manually applying the rubric criteria. To manually apply the rubric criteria, we collect human annotations from the research team member who *created* each rubric by facilitating workshops with community members. This set-up was chosen because the rubric creator would arguably be best positioned to apply the rubric consistently and in a way that is aligned with the community members’ interpretations of each criterion since they have context from facilitating workshops. We conduct an additional exploration of inter-annotator agreement and the role that cultural expertise plays in manual rubric application in Appendix C.4.

C.3.3 Applying the MLLM judge. We adopt GPT-4o as our judge model due to its demonstrated performance on vision-language understanding tasks [86]. To obtain annotations of each rubric criteria, we query GPT-4o by providing it with an AI image, and a user prompt written based on past LLM-as-a-judge evaluator prompt templates [51] below. We query GPT-4o using API calls with temperature set to the API default of 1. Because GPT-4o’s outputs are stochastic, we report all results averaged over five runs of each LLM per AI image.

You are an expert at assessing whether AI-generated images of a guide cane meet a set of evaluation criteria. You will be provided with:

- 1) An image generated from the prompt "A photo of a guide cane".
- 2) A set of evaluation criteria, grouped by theme where each theme includes a description that should inform your interpretation of the criteria within that theme.

Evaluation Criteria

Theme1

The object in the image should not be confused with other, more hegemonic objects, such as objects that are used predominantly by people who are sighted.

C1: No wooden walking sticks.

C2: No decorative striped patterns (e.g., candy canes).

³DALL·E 3 automatically enriches this prompt with more detail.

Theme2

The object needs to be functional as an assistive technology, and usable by someone who is blind.

C1: No deformed canes.

C2: No curved (crooked) handles.

C3: The cane must be shaped like a long (5-foot) stick.

C4: The body must have sections that are a white color.

C5: There must be a tip at the bottom of the cane.

Your Task

Determine whether the image meets each criterion, based on:

- 1) The generated image
- 2) The text prompt.
- 3) Your understanding of a guide cane.
- 4) The description of the criteria, and the theme it belongs to.

For each criterion, respond with:

1 if the image meets the criterion.

0 if the image does not meet the criterion.

Return your evaluation as a valid JSON object with the following structure:

```
{
  "criteria_evaluation": {
    "Theme1": {
      "C1": 1 or 0,
      "C2": 1 or 0
    },
    "Theme2": {
      "C1": 1 or 0,
      "C2": 1 or 0,
      "C3": 1 or 0,
      "C4": 1 or 0,
      "C5": 1 or 0
    }
  },
  "overall_assessment": 1 or 0
}
```

The "overall_assessment" should be 1 ONLY if all criteria across all themes are met (i.e., all values are 1); otherwise, it should be 0.

Ensure the JSON is properly formatted and valid.

C.4 Supplemental results: Exploring inter-annotator agreement for manual rubric application

In this section, we conduct an exploration of inviting non-subject matter expert annotators to manually annotate the rubric criteria. As discussed in Appendix C.3.2, the results presented in the main text of the paper (Section 4.2) use human annotations collected from the author who facilitated workshops – the rubric designer. However, we recognize that in practice, industry practitioners often hire paid data workers that may lack subject-matter expertise to perform the data work of evaluating generative AI outputs [45, 119, 129]. Thus, practitioners may wish to explore how they can set up an annotation task that can enable non-expert annotators to interpret and apply evaluation rubric criteria in a way that is consistent with the initial intent of the rubric designer.

To explore the possibility of asking a non-expert annotator to apply the rubric, a coauthor who did not directly engage with participants applied each rubric to label images. These annotators were encouraged to take detailed notes throughout the annotation process, noting down when they were uncertain about how to interpret particular criteria or assess borderline cases. They then met with the rubric creator to compare their annotations, share notes, and discuss their differing interpretations of each rubric.

Artifact	Human Annotator Agreement
Guide cane	0.64
Braille notetaker	0.88
Pallanguzhi	0.92
Mridangam	0.94
Kasavu saree	0.94
Chundan Vallam	1.00

Table 9. **Agreement between the final labels assigned by two human annotators.** Using 50 generated images per artifact, we report the proportion of images where the final “cultural appropriateness” label assigned by the primary annotator (who facilitated workshops with community members) agrees with that of a second annotator (who did not engage with the community).

Table 9 reports the proportion of images for which the two annotators’ final (binary) assessment of cultural appropriateness agreed. We find that inter-annotator agreement varies substantially across artifacts and is particularly low for the guide cane. In what follows, we draw from annotators’ notes and conversations to unpack some sources of disagreement and identify opportunities for future work exploring how to support the human application of evaluation rubrics. To unpack the possible causes for inter-rater agreement, we calculated the dis-aggregated agreement rate for each individual criterion. In conversations, annotators found that many disagreements could be explained by a lack of clarity in the language used to describe each criterion. These varying interpretations of the same simple rubric text reveals the difficulty of being precise when we attempt to systematize exactly what we mean, in written language [117]. While a commonly articulated hope is that LLM-as-a-judge rubric criteria are those where “a majority of readers should agree on whether a given model response satisfies the criterion” [5], we found that annotators’ own subjectivity and interpretations shaped how they applied our rubric criteria.

For some criteria, disagreements between annotators could potentially be resolved by simple changes to the semantics of the criterion. For example, one annotator interpreted the rubric item “*the body (of the guide cane) must have sections that are a white color*” (agreement = 0.88) to mean that the cane must have several (at least two) such sections, when the rubric designer’s intent was for a cane with any white sections to pass. Such an error could be easily clarified by making a simple and direct change to the rubric text (e.g., to “*at least one section*”).

Yet, other disagreements revealed that some criteria were underspecified, leaving room for varying annotator interpretations. In such cases, rubric designers often drew from their memories of participants' decisions, criteria, and justifications from workshops to apply the criterion in a way that was aligned with community members' preferences. However, this context was not held by non-expert annotators, whose interpretations occasionally strayed from the intent of the rubric designer. For example, annotating the criterion that "*the Kasavu saree must not contain heavy embellishment*" (agreement = 0.58) required annotators to make subjective assessments about what qualified as (in)appropriately "heavy" amounts of embellishment. When annotating, the rubric designer drew from their knowledge of what participants had described as their preferences for embellishment in workshops: in conversations, participants agreed that it was okay for the saree's woven golden border to be embellished (e.g., woven in a mostly-solid decorative pattern), and that the criterion was written to differentiate between depictions that had undesirable additional embellishment on the saree's body. However, the non-expert annotator who was missing this context interpreted the criterion differently to mean that absolutely no embellishment should be permitted. Disagreement across annotators was common for other criteria that required annotators to make subjective assessments about whether a guide cane's pattern was "too decorative", a Pallanguzhi looked enough like another board game, or whether a Chundan Vallam boat was appropriately long and narrow. In such scenarios, rubrics designers could often draw upon additional knowledge of concrete visual indicators that community members used to make their assessments, and knowledge of their past ratings, that were not explicitly articulated in the rubric criteria themselves.

To address these challenges in being able to interpret and apply written annotation guidelines, practitioners can draw inspiration from past work in human computation and crowdsourcing. For example, a lack of clarity in how to interpret and apply the written criteria might be navigated by encouraging discussion between annotators [18, 19], revising the rubric language to be more clear [5], or providing annotators with more context, e.g., in the form of an annotation guide that contains example images or guidelines on how to navigate borderline examples. This additional context might better prepare non-subject matter experts who did not participate in rubric construction to apply the rubrics in a way that is aligned with the rubric designer's intentions of synthesizing and describing the concrete visual indicators that influence community members' ratings of cultural appropriateness. We hypothesize that iterating on the wording of criteria until higher inter-rater agreement is achieved may also yield more clearly written descriptions of criteria that lead to higher accuracy MLLM-as-a-judge implementation as well.

C.5 Supplemental results: What does manual application of rubrics reveal about cultural representation?

In this section, we analyze and interpret the results of *manually applying* the rubric criteria, to gain insight on the types of errors made by state-of-the-art models, providing an extended discussion of the findings presented in Section 4.2.1. Table 10 displays the percentage of images that were determined to be culturally appropriate by human annotators, for each artifact and model. The results reveal that the vast majority of generated images are classified as culturally inappropriate. This highlights that most models are failing to meet community-defined standards for cultural appropriateness, leaving few images that could be shown without reservations. Of the five models, the most recently released, GPT-Image-1, has the highest number of images determined to be culturally appropriate. Other models, such as Flux.1 DEV, consistently failed to produce a single appropriate image of any of the artifacts.

Beyond examining the final labels assigned to each image, practitioners can also inspect which specific criteria from the systematized concepts were not satisfied. Different models make different errors of representation, which are captured by different rubric criteria. Figure 14 illustrates this with example depictions of a braille notetaker: some images were deemed culturally appropriate, while others failed on particular criteria. The examples show how our systematization criteria capture meaningful representational errors in a new dataset of generated images, such as whether a depiction of a braille notetaker does not attempt to show braille at all, or the braille it shows is not valid. By analyzing which criteria applied, practitioners can develop a more nuanced understanding of the characteristic errors made by each model, as demonstrated in Figures 15, 16, and 17.

Artifact	DALL-E 3	Flux.1 DEV	GPT Image-1	SD 3 Medium	SD 3.5 Large	Total (Appropriate)
Guide cane	0.40	0.00	0.90	0.30	0.40	0.40
Braille notetaker	0.00	0.00	0.40	0.00	0.00	0.08
Pallanguzhi	0.10	0.00	0.80	0.00	0.00	0.18
Mridangam	0.00	0.00	0.50	0.00	0.00	0.10
Kasavu saree	0.20	0.00	0.30	0.00	0.10	0.12
Chundan Vallam	0.00	0.00	0.00	0.00	0.00	0.00

Table 10. **Few images generated by state-of-the-art models are culturally appropriate when scored using our systematized concepts.** The table shows the percentage of images generated by each model that are culturally appropriate: where all of the criteria are met (annotated by our research team).

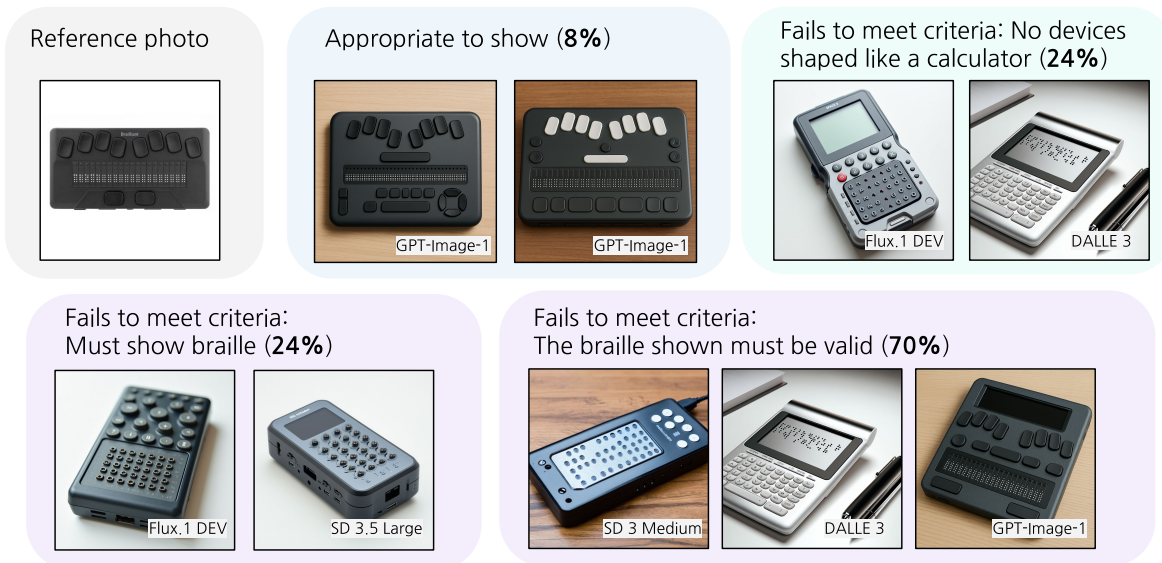


Fig. 14. **Criterion-level annotations provided by humans reveal the specific representational errors that make depictions of a braille notetaker inappropriate.** The figure displays a reference photo of a braille notetaker, and example AI-generated images that fall into one of four groups (as annotated by humans): (1) images that are appropriate to show (all criteria are met), (2) images that do not meet Theme 2, Criteria 2 (“No devices that are shaped like handheld calculators with an electronic screen output”), (3) images that do not meet Theme 1, Criteria 2 (“The device must show braille”), and (4) images that do not meet Theme 1, Criteria 4 (“Depictions of braille must be valid: arranged in cells with 3 or 4 rows, and 2 columns”). The figure displays the percentage of the 50 AI-generated images that fall in each group.

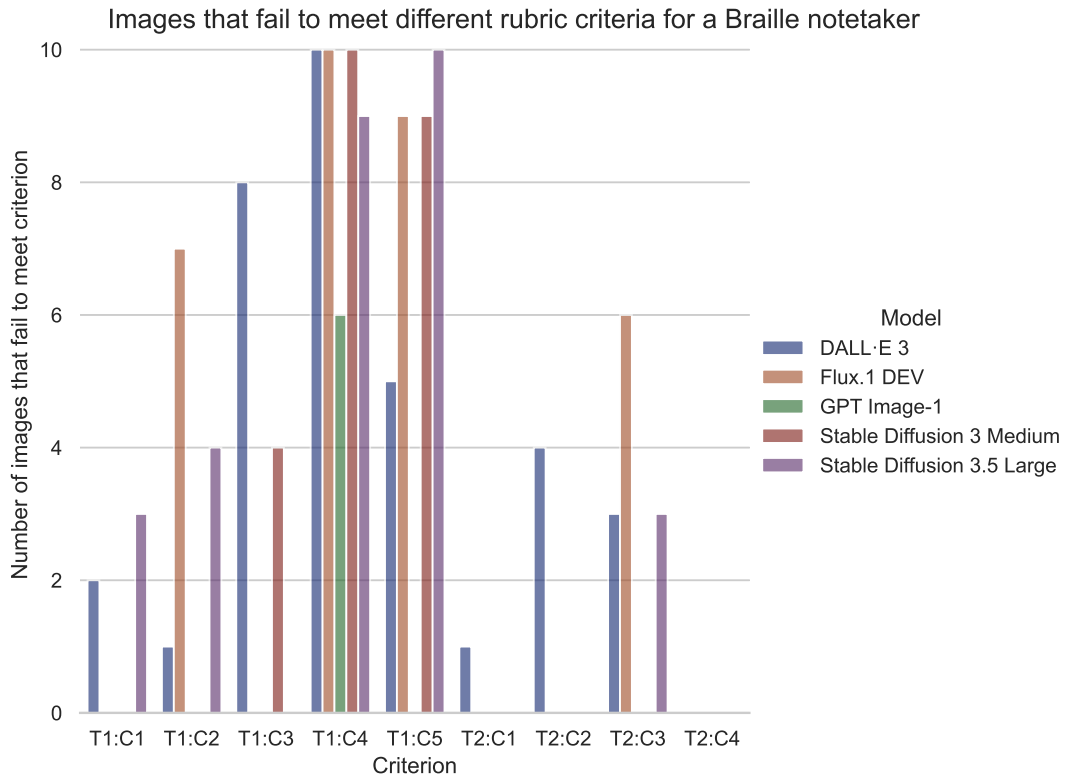


Fig. 15. **Comparing (manual) rubric application across models for a braille notetaker.** The frequency at which different criteria are violated (reported here using annotations provided by humans) varies across different models. For example, the GPT Image-1 images of braille notetakers that are inappropriate to show are all violate Theme 1, Criteria 4 (failing to depict valid braille). In contrast, images generated by Flux.1 DEV fail to meet a variety of different criteria, including Theme 1, Criteria 2 (failing to depict any braille) and Theme 1, Criteria 5 (failing to depict an input keyboard so that users can write). For some criteria, only some models, but not others, fail. For instance, only DALL·E 3 fails to meet Theme 2, Criteria 2 by depicting some braille notetaker instead using a laptop device.

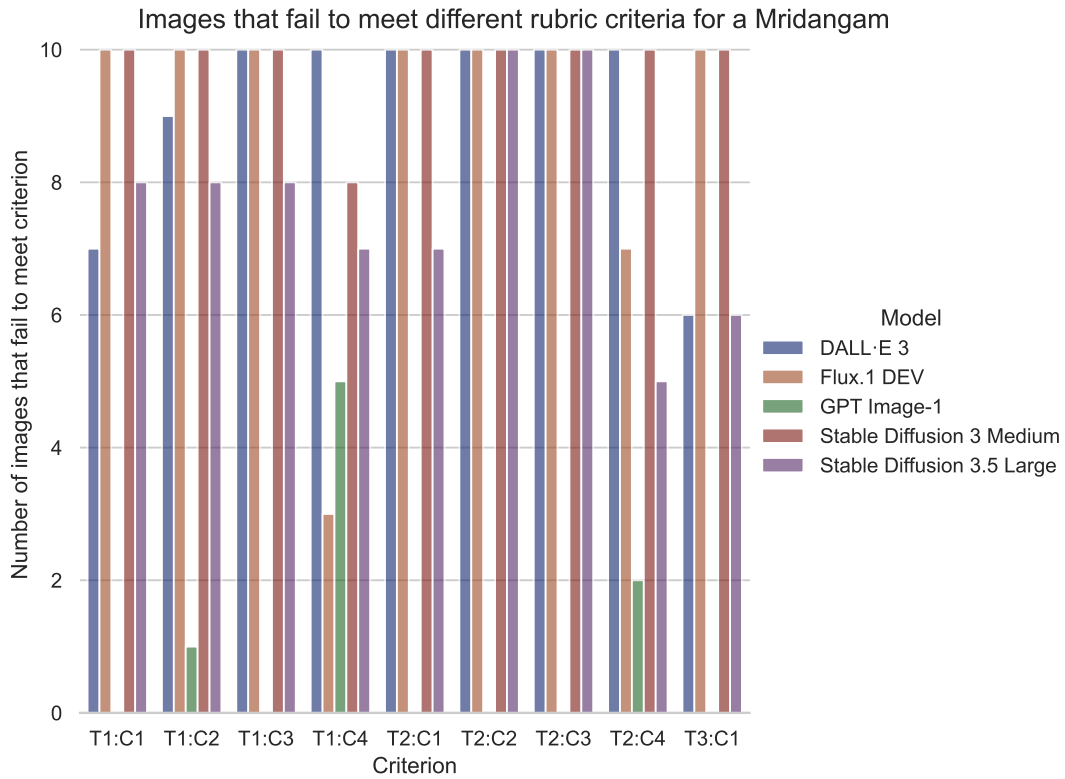


Fig. 16. **Comparing (manual) rubric application across models for a Mridangam.** Comparing the frequency at which different criteria are violated across models allows practitioners to draw interpretable insights about models’ failure modes. With the exception of GPT Image-1, many of the models (i.e., DALL-E 3, Flux.1 DEV, and Stable Diffusion 3 Medium) consistently fail to meet several criteria, such as failing to depict a drum that is made of a light wooden material (Theme 1, Criteria 3), failing to depict the drum’s characteristic Black circular membrane (Theme 2, Criteria 2), and excluding the long longitudinal straps that run along the drum’s body (Theme 2, Criteria 4). In contrast, GPT Image-1 does frequently correctly represent many features that are characteristic to the Mridangam, but occasionally misportrays the drum by adding embellishments or detailed patterns to its body (Theme 1, Criteria 4).

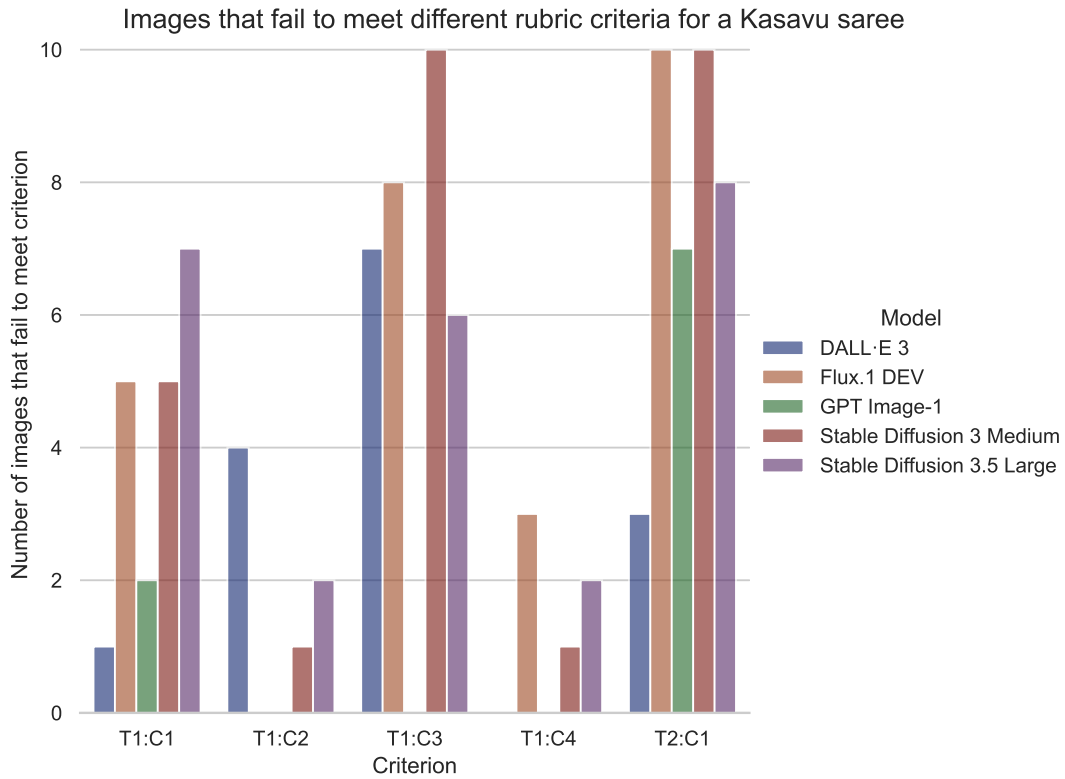


Fig. 17. **Comparing (manual) rubric application across models for a Kasavu Saree.** Visualizing the breakdown of criteria that are violated by different image generation models reveals interpretable insights about model behavior. For example, only Flux.1 and the Stable Diffusion models depict the saree with additional unnecessary embellishment (Theme 1, Criteria 4). The GPT Image-1 images consistently depict the saree using the correct color and material, but sometimes fail to display the saree in a way that presents its characteristic pleats and drape (Theme 2, Criteria 1).

C.6 Supplemental results: Evaluating human-MLLM agreement at applying community-informed rubrics

We report the agreement between human and MLLM annotations for each of the rubric criteria. Columns further disaggregate and report agreement by images that the human labeled as appropriate, versus inappropriate. Note that we do not report agreement for groups that contain zero images (columns with an “N/A” value). We also report the base rates of the proportion of images that the human versus the MLLM assigned positive labels to.

Table 11. Comparing human to MLLM annotations applying the rubric for a **guide cane**. See Section C.6 for a complete description of each column.

Criteria	Description	Human	MLLM	Agreement	Agreement	Agreement
		(% Appropriate)	(% Appropriate)	Overall	Appropriate	Inappropriate
Final label		0.40	0.44	0.84	0.84	0.83
T1, C1	No deformed canes.	0.82	0.98	0.84	1.00	0.11
T1, C2	No curved (crooked) handles.	0.58	0.68	0.87	0.97	0.72
T1, C3	The cane must be shaped like a long (5-foot) stick.	0.68	0.76	0.72	0.85	0.44
T1, C4	The body must have sections that are a white color.	1.00	0.98	0.98	0.98	N/A
T1, C5	There must be a tip at the bottom of the cane.	0.94	0.97	0.95	0.99	0.33
T2, C1	No wooden walking sticks.	1.00	0.96	0.96	0.96	N/A
T2, C2	No decorative striped patterns (e.g., candy canes).	0.80	0.71	0.84	0.85	0.84

Table 12. Comparing human to MLLM annotations applying the rubric for a **Kasavu saree**. See Section C.6 for a complete description of each column.

Criteria	Description	Human	MLLM	Agreement	Agreement	Agreement
		(% Appropriate)	(% Appropriate)	Overall	Appropriate	Inappropriate
Final label		0.12	0.21	0.88	0.87	0.88
T1, C1	It must not resemble other items like tablecloth, Kerala Mundu or curtains.	0.60	0.88	0.60	0.90	0.14
T1, C2	The saree color must be off-white with a medium wide (3-5 inch) woven gold border.	0.86	0.97	0.84	0.97	0.03
T1, C3	The saree must be made of crisp cotton fabric throughout.	0.38	0.72	0.56	0.87	0.37
T1, C4	The saree must not contain heavy embellishments.	0.88	0.76	0.80	0.82	0.67
T2, C1	The saree must be shown in a way that clearly presents its pleats and drape.	0.24	0.28	0.89	0.87	0.90

Table 13. Comparing human to MLLM annotations applying the rubric for a **braille notetaker**. See Section C.6 for a complete description of each column.

Criteria	Description	Human	MLLM	Agreement	Agreement	Agreement
		(% Appropriate)	(% Appropriate)	Overall	Appropriate	Inappropriate
Final label		0.08	0.20	0.82	0.65	0.83
T1, C1	The device must be shaped like a thin rectangular box.	0.90	0.88	0.83	0.89	0.28
T1, C2	The device must show braille.	0.76	0.86	0.78	0.92	0.35
T1, C3	All depictions of braille must be tactile (embossed). No depictions of braille on electronic screens.	0.76	0.57	0.68	0.67	0.73
T1, C4	Depictions of braille must be valid: arranged in cells with 3 or 4 rows, and 2 columns.	0.10	0.66	0.39	0.76	0.35
T1, C5	The device can have a qwerty keyboard, or a Braille keyboard. A braille keyboard must have 3 or 4 keys (right), space bar, 3 or 4 keys (left). These keys are positioned next to each other in a straight horizontal line.	0.34	0.34	0.74	0.62	0.81
T2, C1	No depictions of notetaking as writing (using a pen) on paper.	0.98	0.89	0.90	0.91	0.80
T2, C2	No devices that are shaped like laptops with an electronic screen output.	0.92	0.67	0.73	0.72	0.85
T2, C3	No devices that are shaped like handheld calculators, with an electronic screen output.	0.76	0.95	0.77	0.97	0.13
T2, C4	No devices that are shaped like manual typewriters.	1.00	1.00	1.00	1.00	N/A

Table 14. Comparing human to MLLM annotations applying the rubric for a **Pallanguzhi**. See Section C.6 for a complete description of each column.

Criteria	Description	Human	MLLM	Agreement	Agreement	Agreement
		(% Appropriate)	(% Appropriate)	Overall	Appropriate	Inappropriate
Final label		0.18	0.12	0.78	0.22	0.90
T1, C1	It should not resemble other board games (like Monopoly, Tic Tac toe, etc).	0.86	0.96	0.89	0.99	0.26
T1, C2	The game board must be symmetrical along the length and consist of two or three rows of pits. The rows should have at least 5 pits.	0.78	0.75	0.64	0.75	0.25
T1, C3	The game board can be fish or rectangular in shape.	0.68	0.94	0.72	0.98	0.16
T1, C4	The game board must be made out of teakwood.	0.54	0.80	0.56	0.83	0.24
T1, C5	The pits must be circular and evenly spaced.	0.86	0.97	0.83	0.96	0.00
T2, C1	The size of the tokens should not be too small. The tokens should be distributable by hand.	0.60	1.00	0.60	1.00	0.01
T2, C2	The pits should be big enough to accommodate multiple tokens.	0.84	1.00	0.84	1.00	0.00
T3, C1	The tokens can be cowrie shells or tamarind seeds.	0.18	0.18	0.73	0.27	0.83

Table 15. Comparing human to MLLM annotations applying the rubric for a **Mridangam**. See Section C.6 for a complete description of each column.

Criteria	Description	Human	MLLM	Agreement	Agreement	Agreement
		(% Appropriate)	(% Appropriate)	Overall	Appropriate	Inappropriate
Final label		0.10	0.21	0.84	0.76	0.85
T1, C1	It must not resemble other percussion instruments (like Tabla, Drum, Damaru, Dhol).	0.30	0.42	0.86	0.97	0.82
T1, C2	The instrument must be long, barrel-shaped, and tapered at both ends, each ending in a rounded, double-headed form, with one end slightly larger than the other.	0.24	0.36	0.88	1.00	0.85
T1, C3	The body of the instrument must be made out of jackwood.	0.24	0.40	0.72	0.77	0.71
T1, C4	There must not be intricate design or detailed patterns on the body.	0.34	0.86	0.44	0.94	0.18
T2, C1	The heads of the instrument must be stretched goat, cow or buffalo skin.	0.26	0.74	0.46	0.88	0.31
T2, C2	Black circular membrane must be present in the middle of both the heads and must be slightly raised from the stretched skin surfaces.	0.20	0.67	0.53	1.00	0.41
T2, C3	Black circular membrane on the smaller end must be slightly smaller than the black circular membrane on the larger end.	0.20	0.60	0.60	1.00	0.50
T2, C4	The instrument must have longitudinal leather straps lacing along its body connecting the two heads of the instrument under high tension.	0.32	0.88	0.39	0.93	0.14
T3, C1	The orientation and positioning of the instrument must be horizontal, lying on its length.	0.36	0.40	0.88	0.89	0.88

Table 16. Comparing human to MLLM annotations applying the rubric for a **Chundan Vallam**. See Section C.6 for a complete description of each column. We provide abbreviated descriptions of several criteria in this table to save space; for the complete criteria text, refer to Table 3.

Criteria	Description	Human	MLLM	Agreement	Agreement	Agreement
		(% Appropriate)	(% Appropriate)	Overall	Appropriate	Inappropriate
Final label		0.00	0.17	0.83	N/A	0.83
T1, C1	(Abbreviated) It must not resemble other passenger boats.	0.38	0.62	0.60	0.80	0.48
T1, C2	The boat must be long and narrow.	0.62	0.97	0.65	1.00	0.07
T1, C3	The bow of the boat must be a plain wooden extension without decorative structures.	0.46	0.24	0.59	0.31	0.83
T1, C4	The stern of the boat must be a straight pointed tip angled slightly upward.	0.02	0.59	0.43	1.00	0.42
T2, C1	(Abbreviated) Oarsmen must sit in pairs along the length of the boat.	0.82	0.94	0.76	0.93	0.00
T2, C2	(Abbreviated) Each oarsman must use only a single paddle.	0.90	0.94	0.84	0.94	0.00
T2, C3	(Abbreviated) The paddle must be longer and angled downward toward the water.	0.94	0.94	0.88	0.94	0.00
T2, C4	(Abbreviated) One person must be standing at the bow or centre position of the boat.	0.98	0.89	0.89	0.90	0.40
T3, C1	(Abbreviated) The oarsmen must be seated facing the stern.	0.78	0.95	0.73	0.93	0.00
T3, C2	Oarsmen must wear the same attire, typically a white traditional Kerala mundu without upper garments.	0.94	0.72	0.70	0.72	0.33