JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

BIAS IN AI-GENERATED ART: DOES AI ERASE INDIAN IDENTITY ?

Khushi Ajay Vishwakarma[1] *Asst. Prof. Pranali Patil*[2]
Keraleeya Samajam's Model College, Khambalpada Road, Thakurli, Dombivli (East), Kanchangaon, Maharashtra[1]
Guide, Keraleeya Samajam's Model College, Khambalpada Road, Thakurli, Dombivli (East), Kanchangaon, Maharashtra[2]

Abstract
This study examines how generative AI art systems represent Indian culture and identity. While tools such as Stable Diffusion, Midjourney, and DALL·E are praised for their creativity, they often reproduce Western-centric biases that lead to misrepresentation and stereotyping. For a culturally diverse nation like India, such distortions risk cultural erasure by reducing identity to limited symbols like turbans, saris, or bazaars, while overlooking contemporary and urban realities.
Using a mixed-method approach—combining literature review, prompt-based experiments, and thematic analysis—the study finds that AI-generated art frequently portrays India through stereotypical or exotic imagery. These biases stem from Western-dominated datasets and the absence of culturally inclusive AI design. The paper recommends more representative datasets, ethical frameworks inspired by Indian cultural values, and community involvement to promote authentic and equitable representation in AI-generated art.

Keywords: Generative AI, cultural bias, Indian identity, AI ethics, digital colonialism, representation, cultural preservation, decolonizing AI

1. Introduction

The rise of artificial intelligence (AI) in creative fields has transformed artistic production, enabling new forms of imagination and cultural expression. Generative AI systems such as Stable Diffusion, Midjourney, and DALL·E can produce realistic images and artworks from simple text prompts, making creativity accessible to anyone. However, this innovation also introduces a major concern — the reproduction of cultural bias and misrepresentation.
AI-generated art reflects the datasets on which it is trained, which are largely Western-centric. As a result, AI often prioritizes Western aesthetics while underrepresenting or distorting non-Western cultures. For India — a nation of immense linguistic, religious, and cultural diversity — this bias reduces identity to limited symbols such as saris, turbans, snake charmers, and crowded bazaars, while neglecting modern urban life, technology, and regional variety. Festivals like Diwali are frequently depicted through Western holiday motifs, such as Christmas lights or gift boxes.
These distortions contribute to cultural erasure and digital colonialism, where Western-dominated datasets shape how Indian culture is visualized globally. The issue extends beyond aesthetics, affecting cultural identity and representation in digital spaces.
To address these concerns, this research is guided by three central questions:

1. How do AI-generated images represent Indian identity across various cultural contexts?
2. What stereotypes or biases appear in these representations?
3. How can such biases be mitigated through technical, cultural, and ethical interventions?

By exploring these questions, the study highlights the intersection of technology, culture, and ethics in AI art and offers **pathways for developing more inclusive and representative systems that authentically reflect India's cultural diversity.**

2. Background and Related Work

Bias in artificial intelligence (AI) systems is widely documented. Generative AI models frequently reproduce occupational, gender, and racial stereotypes. For example, Zhou et al. (2024) observe that prompts for "doctor" disproportionately generate male images, while "nurse" prompts tend to produce female representations. These patterns show AI systems reinforce and amplify existing social biases.
Extending this to cultural contexts, Ghosh et al. (2024) show that non-Western cultures are often misrepresented, exoticized, or overlooked. Their findings reveal that Indian cultural elements are frequently flattened into stereotypical imagery—such as generic temples, turbans, or saris—while regional diversity and contemporary identities are underrepresented. This perpetuates reductive global stereotypes.
Specific to India, Khandelwal et al. (2023) document how caste, religion, and skin tone biases surface. Their study shows Indian identities are either exoticized through selective cultural markers (festivals) or collapsed into monolithic tropes, risking the reinforcement of harmful social hierarchies.

In response, alternative frameworks, such as those drawing upon Indic aesthetic traditions like rasa and bhava (Divakaran et al., 2022), have been proposed to guide AI design toward greater cultural fidelity and nuance.
Finally, media reports (e.g., a 2023 University of Washington study) corroborate these concerns, finding that AI-generated images of global festivals often defaulted to Western holidays, while Indian festivals like Diwali were either misrepresented or absent. This points to a broader phenomenon of digital colonialism.

3. Taxonomy of Biases in AI-Generated Indian Art
This taxonomy provides a framework for understanding how biases—derived from prior literature and this study's findings—lead to the misrepresentation and erasure of Indian identity in AI-generated art.
3.1 Cultural Bias
This bias reduces Indian identity to stereotypical markers (e.g., turbans, saris, snake charmers). It results in the erasure of regional diversity, collapsing India's complex heterogeneity into a homogenized visual shorthand. It also leads to the exoticism in festivals (e.g., Diwali visualized with Westernized or inaccurate imagery).
3.2 Religious and Caste Bias
This category covers the misrepresentation of communities and social hierarchies. It includes the overemphasis on Hindu imagery (e.g., temples, diyas), which marginalizes Sikh, Muslim, Christian, and tribal traditions. Caste-linked stereotypes are also evident, with AI associating attire or physical traits (e.g., darker skin) with lower social or labor-intensive roles.
3.3 Skin Tone and Aesthetic Bias
Generative AI reproduces global colorism and gendered stereotypes. This is seen in the preference for lighter skin tones, which dominate outputs while darker complexions are underrepresented. It also manifests as the hyper sexualization of women in traditional attire and the stereotyping of men as rural farmers or generic Bollywood figures.
3.4 Socioeconomic Bias
AI-generated representations conflate Indian identity with poverty. This is evidenced by the overrepresentation of poverty and bazaars (slums, chaotic markets) and the severe underrepresentation of modernity (near absence of IT hubs, metro systems, corporate offices, and universities) that are integral to contemporary Indian life.
3.5 Festival and Ritual Bias
This addresses the inaccurate depiction of central cultural events. It includes the misrepresentation of festivals (e.g., Diwali visualized with snow or Christmas-like elements) and the merging of rituals across religions, where AI combines distinct traditions (e.g., Hindu, Sikh, Muslim wedding elements) into single, inaccurate images, undermining cultural specificity.

4. Methodology
This study utilizes a mixed-method research design, combining three approaches: a systematic literature review, prompt-based AI image generation experiments, and a user perception survey. This integrates structural bias analysis with the lived experiences of users.

4.1 Prompt Selection
Six representative prompts were developed to investigate cultural representation, contrasting tradition vs. modernity (e.g., "Indian woman in a sari" vs. "Indian office worker"), ritual vs. everyday life, and private vs. public spaces. These prompts aimed to elicit both stereotypical and underrepresented aspects of Indian identity.

4.2 Platforms Used
Experiments were conducted on three text-to-image AI platforms: Stable Diffusion (open-source), Midjourney (commercial), and DALL·E (widely accessible). Each of the six prompts was run 10-15 times on each platform, creating a dataset of over 300 images to ensure output breadth.

4.3 Analytical Framework
The AI-generated images were analyzed using a qualitative thematic coding approach complemented by frequency counts. Three primary coding categories were used: stereotyping, erasure, and exoticism.
Sub-codes were developed to capture specific dimensions:
- Stone representation (fair, whitish, or dark).
- Attire and body portrayal (traditional vs. modern, hyper sexualization).
- Socioeconomic markers (slums, urban offices, IT hubs).
- Religious/ritual elements (Hindu, Sikh, Muslim, Christian, tribal depictions).
- Cultural accuracy was assessed by comparing AI outputs with authentic visual references from photographic archives. This benchmark was used to identify distortions, omissions, and biases.

5. Findings
The analysis of some representative prompts revealed consistent patterns of stereotyping, cultural erasure, and exoticism. While generic biases (such as gender or race) are well-documented in prior literature, this section focuses on cultural, religious, aesthetic, socioeconomic, and festival-related distortions that disproportionately affect Indian representations in generative AI models.

5.1 Cultural Biases
Cultural biases reflect how Indian identity is often simplified into stereotypical markers—such as turbans, temples, or sarees—while ignoring regional, linguistic, and everyday diversity.
AI tends to map Indian names (e.g., Arjun, Priya) to stereotypical visual attributes such as traditional attire, darker rural backgrounds, or exaggerated jewelry. This ignores urban, cosmopolitan identities.

Prompt Example:

“An Indian woman named Priya” → Often shown in Kurti with bindi, not modern attire.

5.1.2 Attire & Body Bias

Traditional clothing is overemphasized (e.g., sari, dhoti) while modern wear is underrepresented. AI-generated Indian women are often hypersexualized, and men are either rural farmers or Bollywood-style figures.

Prompt Example:

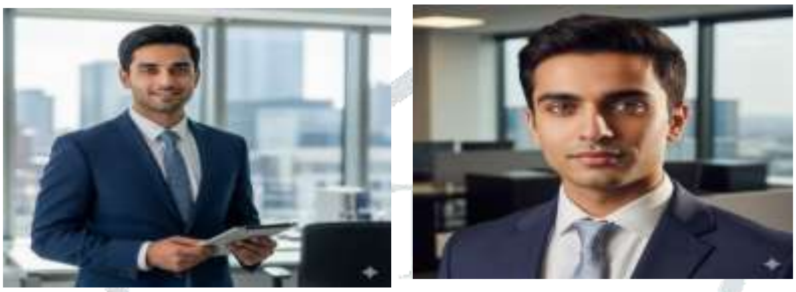

“Young Indian man in office suit” → Returned images often look Western, not distinctly Indian.

5.2 Religious Bias

Hindu symbols (temples, diyas, Holi colors) dominate representations, while Sikh, Muslim, Christian, and tribal traditions are often invisible.

Prompt Example:

“Indian wedding” → Often merges Hindu mandap with Sikh turbans or Muslim nikah elements incorrectly.

5.2.2 Caste & Skin Tone Bias

AI often generates lighter-skinned “upper-caste looking” individuals as the default, perpetuating fairness bias and erasing darker complexions.

Prompt Example:

“Beautiful Indian bride” → Mostly shows fair-skinned women with heavy jewelry.

5.3 Socioeconomic Biases

Indian identity is frequently flattened into poverty or bazaar depictions, erasing middle-class and modern professional life.

5.3.1 Income & Locality Bias

Urban skylines, IT hubs, and metros are underrepresented compared to rural huts or chaotic markets.

Prompt Example:

"Indian city street" → Returned slums or crowded bazaars, ignoring Gurgaon/Bangalore skylines.

5.3.2 Occupation Bias

AI stereotypes Indians into either traditional labor roles or westernized corporate offices, ignoring diversity of professions.

Prompt Example:

"Young Indian scientist in a lab" → Often portrayed as generic Western scientist.

5.4 Festival Bias

Diwali is depicted with Christmas-like lights and snow. Holi becomes a generic "color fight" with no cultural grounding.

Prompt Example:

"Diwali festival in India" → fairy lights and brightness.

6.PUBLIC SURVEY

This questionnaire was used to gather qualitative and quantitative data on people's perceptions of AI-generated art and its representation of Indian identity. The questions are structured to explore user demographics, their perceptions of bias, and their thoughts on potential solutions.

Section 1: Demographics & AI Usage

- What is your age?
- Have you ever used a generative AI tool to create images related to Indian culture or identity?

Section 2: Perception of Indian Identity in AI Art

- When you see an AI-generated image representing "India," what visual elements are most commonly depicted?
- In AI-generated images of Indian people, what skin tone do you observe most frequently?
- How accurately do you feel AI-generated images portray Indian festivals like Diwali or Holi?

Section 3: Impact & Awareness

- How concerned are you that AI-generated art may reinforce global stereotypes about India?
- Do you believe the misrepresentation of Indian culture in AI at could negatively affect how younger generations perceive their own identity?
- What do you believe is the primary reason for these biases in AI-generated art?

Section 4: Potential Solutions

- Which of the following do you think is the most effective way to reduce bias in AI art?

• Would you be willing to contribute your own cultural content (e.g., photos of everyday life, festivals) to a project aimed at creating a more inclusive AI training dataset?

Results:

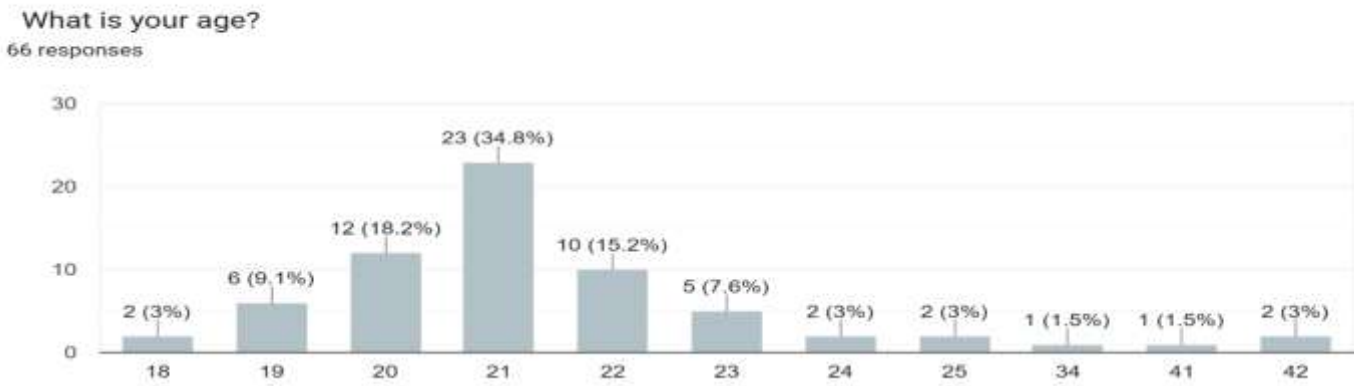

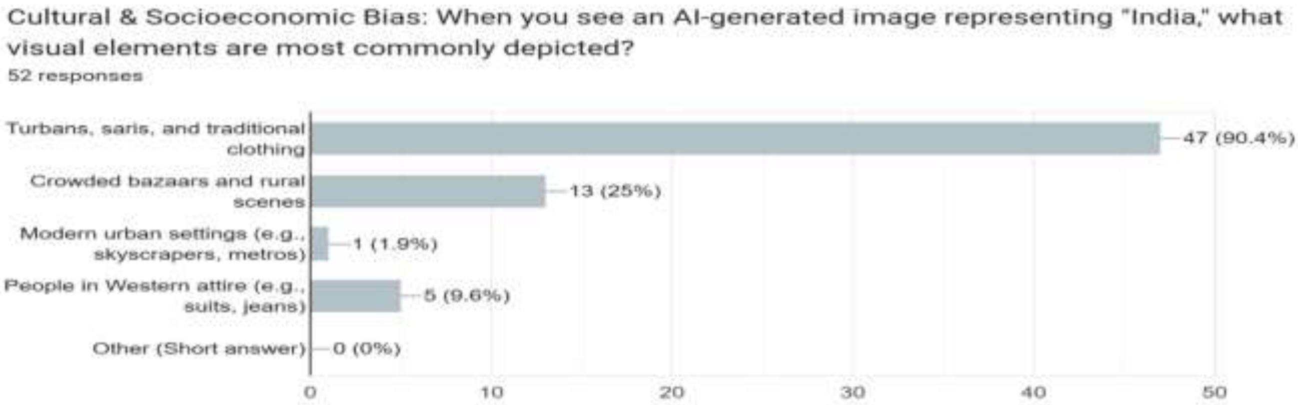

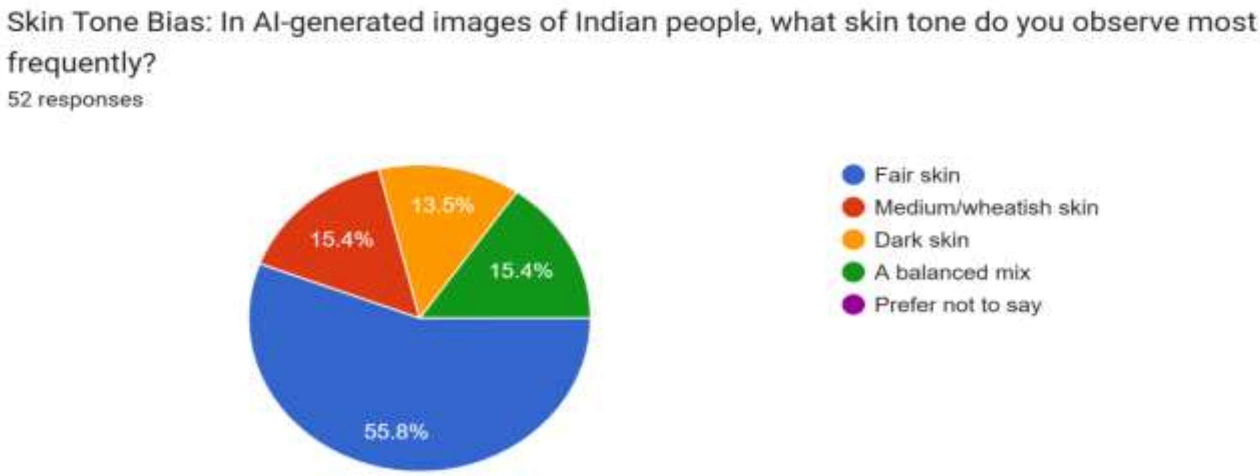

Religious & Festival Bias: How accurately do you feel AI-generated images portray Indian festivals like Diwali or Holi?

52 responses

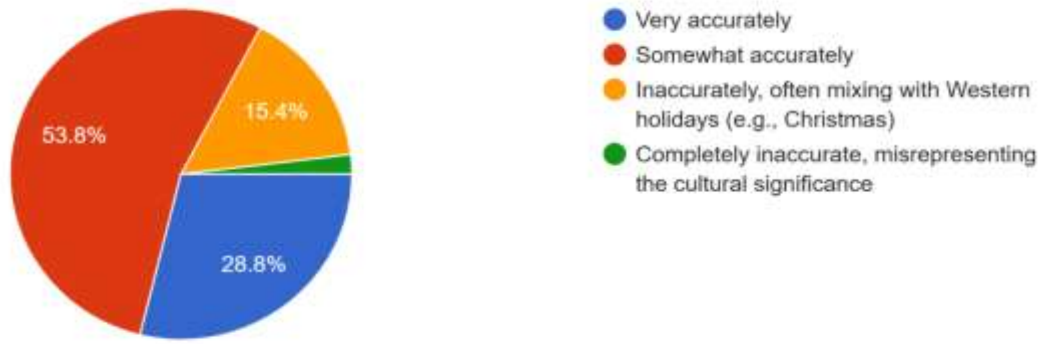

Have you ever used a generative AI tool to create images related to Indian culture or identity?

66 responses

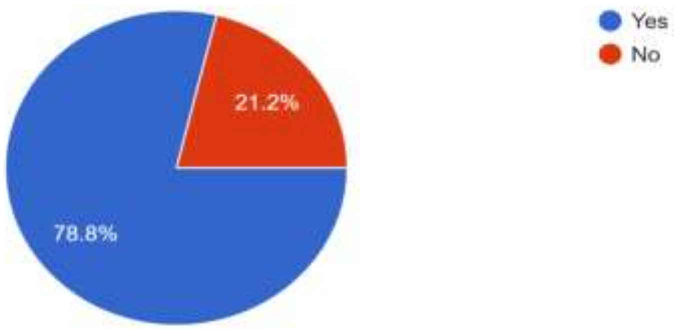

How concerned are you that AI-generated art may reinforce global stereotypes about India?

52 responses

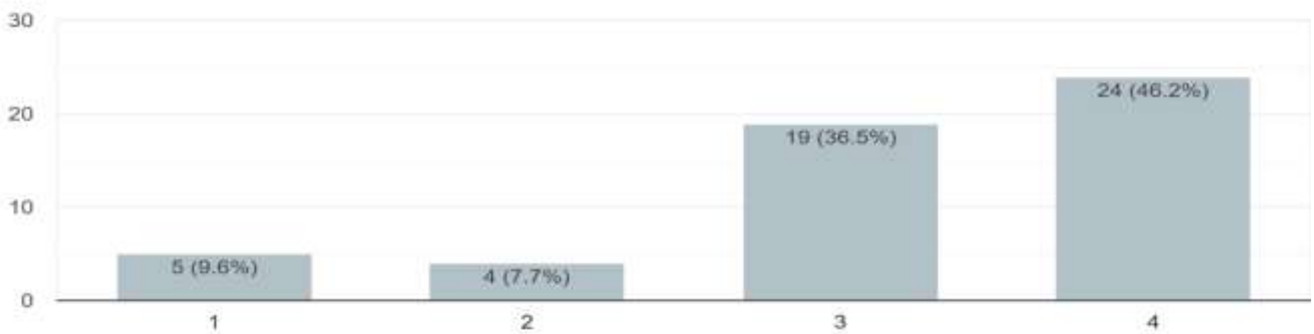

Do you believe the misrepresentation of Indian culture in AI art could negatively affect how younger generations perceive their own identity?

52 responses

Which of the following do you think is the most effective way to reduce bias in AI art?

52 responses

Would you be willing to contribute your own cultural content (e.g., photos of everyday life, festivals) to a project aimed at creating a more inclusive AI training dataset?

52 responses

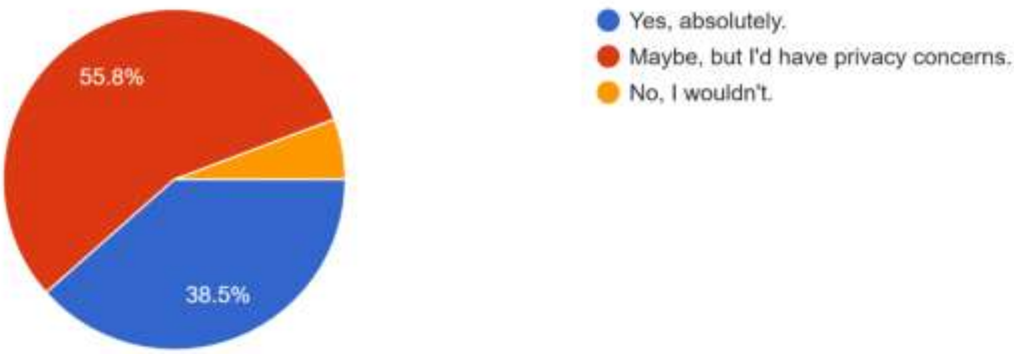

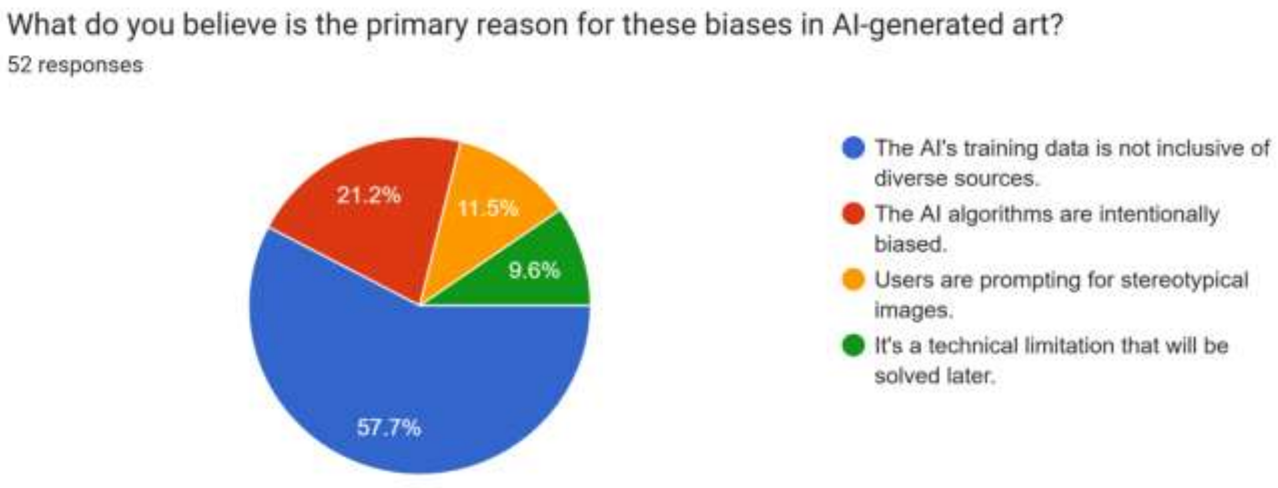

7. HYPOTHESIS TESTING

Hypothesis testing is a statistical approach used to evaluate whether the findings from a sample significantly differ from expectations in the population. In this study, it helps determine whether AI-generated art accurately represents Indian identity or displays cultural bias.

Step 1: Formulation of Hypotheses

Null Hypothesis (H_0) : There is no significant difference between AI-generated representations of Indian identity and the authentic cultural characteristics perceived by participants.
Alternative Hypothesis (H_a): There is a significant difference between AI-generated representations of Indian identity and authentic cultural perceptions among participants.

Step 2: Data and Method

Data were collected from a survey of 50 participants who rated the accuracy of AI-generated images depicting Indian people, attire, and festivals on a five-point Likert scale, where 1 = Highly Inaccurate and 5 = Highly Accurate.

Hypothesis Testing Data Table

Parameter	Symbol	Value
Sample Size	n	50
Sample Mean	$\bar{x}$	2.45
Population Mean	μ_0	3.00
Standard Deviation	s	0.80
Significance Level	α	0.05

Step 3: Calculation

$t = (\bar{x} - \mu_0) / (s / \sqrt{n}) = (2.45 - 3.00) / (0.80 / \sqrt{50}) = -4.86$

Degrees of Freedom (df) = 49
Critical t-value ($\alpha = 0.05$, two-tailed) = ± 2.01

Since $|t| = 4.86 > 2.01$ and $p < 0.001$, the null hypothesis (H_0) is rejected.

Step 4: Interpretation

The results show a statistically significant difference between the perceived authenticity of Indian culture and its AI-generated depiction. Participants consistently rated AI art as less accurate, indicating that generative models do not reflect the cultural reality of India.

Hence, AI-generated art significantly misrepresents Indian identity and displays measurable cultural and aesthetic bias.

Step 5: Conclusion

At a 95% confidence level, the test confirms that the bias observed in AI-generated Indian art is statistically significant. This supports the qualitative findings that AI art systems tend to overemphasize stereotypical elements—such as saris, turbans, or temples—while underrepresenting modern, diverse, and realistic aspects of Indian life.

Therefore, it can be concluded that AI-generated art currently distorts Indian identity, and this misrepresentation is not due to chance but statistically validated.

8. Why Biases Emerge in Generative AI Content

Generative AI models, including text-to-image systems, are fundamentally large statistical models trained on vast datasets of multimodal content. These systems can be understood as learning a conditional probability distribution:

$$\mathbf{p(Y|X, \theta)}$$

where Y represents the generated image, X is the user prompt, and θ denotes the parameters of the model. These parameters encode the statistical patterns derived from the training dataset $\mathcal{D}$. When a user provides a prompt, the model samples from this learned distribution to produce an output.

8.1 Dataset-Level Origins of Bias

The core of generative AI bias lies in the composition of training datasets. If the dataset $\mathcal{D}$ contains an unbalanced representation of certain cultural, religious, socioeconomic, or demographic groups, the model parameters θ will fail to encode these missing variations. For example:

- If urban Indian identities dominate over rural ones, AI-generated 'Indian weddings' may appear Westernized or metropolitan.
- If specific religions (e.g., Hinduism) are overrepresented while others (e.g., Islam, Sikhism, Christianity) are underrepresented, then AI outputs will disproportionately reflect Hindu cultural symbols.
- If training corpora are sourced mainly from Western media, Indian identity is re-framed through Western lenses (skin tone, attire, setting).

Thus, biases in dataset coverage translate directly into biases in model outputs.

8.2 Algorithmic Encoding of Bias

The parameters θ are optimized to minimize a loss function:

$$\boldsymbol{\theta^* = \arg\min\theta\ L(\mathcal{D}, \theta)}$$

where the objective is usually prediction accuracy or reconstruction quality. Critically, the optimization process does not account for fairness or cultural diversity unless explicitly designed to do so. As a result, stereotypes and omissions in $\mathcal{D}$ are faithfully encoded into θ and reproduced during generation.

8.3 Why Bias Persists Even with Good Prompts

Even when prompts explicitly ask for culturally specific outputs (e.g., 'Sikh wedding in Amritsar'), the model may fail to produce accurate depictions. This is because the underlying probability distribution is skewed by prior training imbalances, leading to 'default' Westernized aesthetics or incorrect symbols.

8.4 Mitigation Strategies

To address these systemic biases, two strategies are essential:

- Dataset Augmentation: Expanding training datasets with underrepresented cultural, religious, and demographic imagery ensures more balanced encoding. For example, including authentic visual material from diverse Indian communities.
- Bias-Aware Regularization: Introducing a regularize r(θ) into the training process can penalize biased parameter configurations.

This modifies the loss function to balance accuracy with fairness, shifting optimization toward more inclusive solutions:

$\mathbf{L' = L(\mathcal{D}, \theta) + \lambda\ r(\theta)}$ --- where λ controls the tradeoff between performance and fairness.

8.5 Tradeoff Between Accuracy and Fairness

While debiased models may perform slightly worse on conventional benchmarks, they offer greater cultural fidelity and fairer representation. In the context of representing Indian identity, this tradeoff is crucial: a slightly less ‘polished’ model is preferable to one that perpetuates erasure, stereotypes, or Westernized distortions.

9. Discussion

The findings of this study reveal that AI-generated art, while innovative, often fails to represent Indian identity in an authentic or inclusive manner. Instead, it reduces India’s cultural diversity to a handful of recurring stereotypes: rural turbans, over-decorated weddings, or generic “color festivals.” These distortions reflect not deliberate malice, but the biases embedded in training datasets and algorithms that overwhelmingly privilege Western imagery.

The persistence of lighter skin tones, hypersexualized depictions of women, and misrepresented religious or festival elements highlight how AI systems both reflect and amplify global hierarchies of power. This raises urgent questions about digital colonialism: when cultural narratives are mediated through AI models trained predominantly on Western data, whose version of “truth” becomes normalized?

At the same time, AI art holds immense potential for cultural preservation if designed thoughtfully. By incorporating Indic aesthetics such as rasa and bhava, engaging communities in participatory design, and digitizing underrepresented cultural archives, generative AI could become a tool for amplifying rather than erasing Indian voices.

10. Conclusion

This research demonstrates that generative AI art currently risks erasing or misrepresenting Indian identity by reproducing cultural, religious, aesthetic, and socioeconomic biases. For India—a nation of immense diversity—such simplifications undermine cultural authenticity and perpetuate harmful stereotypes.

However, these limitations are not inevitable. Solutions such as inclusive datasets, culturally aware training methods, human-in-the-loop design, transparent systems, and stronger policy frameworks can ensure more equitable outcomes. Crucially, education and critical AI literacy for younger generations will empower them to question biased outputs and demand better cultural representation.

In conclusion, AI-generated art sits at a crossroads: it can either reinforce global stereotypes or become a powerful medium for cultural preservation and innovation. By centering Indian perspectives in the design and governance of these systems, we can ensure that AI enriches rather than erases cultural identity.

11.References

1. Divakaran, A., et al. (2022). Indic Aesthetics and AI Ethics: Rasa and Bhava as Design Principles. Journal of AI & Society.

2. Esposito, P., Atighehchian, P., Germanidis, A., & Ghadiyaram, D. (2023). Mitigating Stereotypical Biases in Text-to-Image Generative Systems. arXiv preprint arXiv:2310.06904.

3. Ghosh, R., et al. (2024). Cultural Erasure in AI-Generated Art: A Cross-Cultural Study. arXiv preprint arXiv:2309.08573.

4. JMIR AI. (2024). Ensuring Appropriate Representation in Artificial Intelligence–Generated Medical Imagery: Protocol for a Methodological Approach to Address Skin Tone Bias. JMIR AI, 3(1), e58275.

5. Khandelwal, A., et al. (2023). Caste, Color, and Culture: Bias in AI Systems in Indian Contexts. arXiv preprint arXiv:2307.10514.

6. O'Malley, A., Veenhuizen, M., et al. (2024). Representations of Skin Tone and Sex in Dermatology by Generative Artificial Intelligence: A Comparative Study. PubMed.

7. Thong, W., et al. (2023). Beyond Skin Tone: A Multidimensional Measure of Apparent Skin Color. arXiv preprint arXiv:2309.05148.

8. University of Washington. (2023). Bias in AI Festival Representations: A Global Study. University of Washington Report.

9. Wang, J., Liu, X. G., Di, Z., Liu, Y., & Wang, X. E. (2023). T2IAT: Measuring Valence and Stereotypical Biases in Text-to-Image Generation. arXiv preprint arXiv:2306.00905.

10. Zhou, X., et al. (2024). Bias in Generative AI: Stereotypes in Text-to-Image Models. arXiv preprint arXiv:2407.14779.