

SANSKRITI: A Comprehensive Benchmark for Evaluating Language Models’ Knowledge of Indian Culture

Arijit Maji¹, Raghvendra Kumar¹, Akash Ghosh¹, Anushka², Sriparna Saha¹

¹Department of Computer Science and Engineering, Indian Institute of Technology Patna, India

²Department of Computer Science, Banasthali Vidyapeeth University, Rajasthan, India

{arijit_2311ai25, raghvendra_2221cs27, akash_2321cs19, sriparna}@iitp.ac.in, guptaanushka024@gmail.com

Abstract

Language models (LMs) are indispensable tools shaping modern workflows, but their global effectiveness depends on understanding local socio-cultural contexts. To address this, we introduce *SANSKRITI*, a benchmark designed to evaluate language models’ comprehension of India’s rich cultural diversity. Comprising of 21,853 meticulously curated question-answer pairs spanning 28 states and 8 union territories, *SANSKRITI* is the largest dataset for testing Indian cultural knowledge. It covers sixteen key attributes of Indian culture namely rituals and ceremonies, history, tourism, cuisine, dance and music, costume, language, art, festivals, religion, medicine, transport, sports, nightlife and personalities, providing a comprehensive representation of India’s cultural tapestry. We evaluate *SANSKRITI* on leading Large Language Models (LLMs), Indic Language Models (ILMs), and Small Language Models (SLMs), revealing significant disparities in their ability to handle culturally nuanced queries, with many models struggling in region-specific contexts. By offering an extensive, culturally rich, and diverse dataset, *SANSKRITI* sets a new standard for assessing and improving the cultural understanding of LMs.

1 Introduction

Language Models (LMs) have revolutionized the way we interact with technology, becoming indispensable tools in applications such as natural language understanding, content generation, knowledge discovery, and decision-making (Brown et al., 2020; Ghosh et al., 2024c; Devlin et al., 2019; Ghosh et al., 2024b). Models like Large Language Models (LLMs), Indic Language Models (ILMs), and Small Language Models (SLMs) are increasingly integrated into workflows across industries, enabling seamless communication and

efficient problem-solving (Ouyang et al., 2022)¹. These advancements, however, come with the critical challenge of ensuring that LMs cater to diverse populations by recognizing and reasoning about their linguistic and cultural contexts (Bender et al., 2021; Ghosh et al., 2025).

While LMs excel in syntactic and semantic language understanding, their effectiveness depends significantly on their ability to comprehend cultural diversity (Lin et al., 2021). A model’s capacity to recognize cultural contexts is crucial for its application in education, governance, healthcare, and entertainment in culturally rich societies (Liang et al., 2022; Ghosh et al., 2024a). The inability to incorporate cultural nuances risks perpetuating biases, stereotypes, or inaccuracies, thereby alienating underrepresented communities. Models that can capture such diversity not only perform better but also foster inclusivity and equity in AI-driven solutions (Blodgett et al., 2020).

India’s unparalleled cultural diversity makes it an ideal testbed for evaluating the cultural competence of LMs. The country is home to 28 states and 8 union territories, each characterized by distinct traditions, festivals, cuisines, music, dance forms, architectural styles, and historical narratives (Sharma, 2019). For example, states like Rajasthan and Tamil Nadu exhibit unique artistic and culinary traditions, while festivals such as Diwali, Eid, and Onam highlight diverse religious and regional practices (Narayanan, 2020; Tripathi, 2019).

This cultural richness presents significant challenges for LMs. Existing benchmarks like TyDi QA (Joshi et al., 2020) and XQUAD (Artetxe et al., 2020) primarily focus on multilingualism, overlooking cultural depth. Models trained on global

¹In this study we define Large Language Model(LLM) as any language model of size more than or equal to 7 billion parameters, Small language models(SLM) are any that has less or equal to 5 billion parameters and Indic Language models (ILM) are any model that is developed inside India indigenously

datasets often fail to capture region-specific nuances, such as the significance of Kathakali performances in Kerala or Madhubani art in Bihar (Patel et al., 2023). This gap in understanding not only limits the utility of these models in culturally complex regions but also hinders their acceptance by local populations (Liang et al., 2022).

Motivation for This Research : The absence of culturally refined benchmarks, especially based on different nuanced attributes, inspired the creation of *SANSKRITI*, a dataset designed to evaluate the cultural understanding of LMs in the Indian context. Unlike existing datasets like Seth et al. (2024a), which are smaller and do not include all Indian states and union territories, *SANSKRITI* focuses on cultural aspects across all 28 states and 8 union territories of India, providing a complete picture of the country’s rich cultural diversity. The motivation behind this research lies in bridging the gap between technological advancements in LMs and their real-world applicability in culturally diverse societies. We developed the *SANSKRITI* dataset to benchmark cultural understanding in India, utilizing diverse data sources to cover key attributes like rituals and ceremonies, history, tourism, cuisine, dance and music, costume, language, art, festivals, religion, medicine, transport, sports, nightlife, and personalities. A team of 40 annotators curated 21,853 questions across four categories: association, country prediction, general awareness, and state prediction.

Research Questions: This research aims to address the following questions:

1. How do various language models (LLMs, SLMs, and ILMs) perform on the *SANSKRITI* dataset, and what are the performance trends across these model categories?
2. How do language models perform across different cultural attributes represented in the *SANSKRITI* dataset?
3. What are the performance trends of language models across different Indian states and union territories?
4. How do language models perform across various question types in the *SANSKRITI* dataset?

Contributions: Our key contributions are summarized as follows:

1. **Introduction of *SANSKRITI*:** We present *SANSKRITI*, a dataset of 21853 manually curated and validated question-answer pairs, designed to evaluate the cultural competence of LMs in the

Indian context.

2. **Comprehensive Cultural Coverage:** *SANSKRITI* encompasses 16 key cultural attributes spanning all Indian states and union territories.

3. **Benchmarking Across Models:** We evaluate leading LLMs, ILMs, and SLMs on *SANSKRITI*, identifying critical gaps in their ability to reason about culturally nuanced queries.

4. **Public Availability:** The *SANSKRITI* dataset and benchmarking results are made publicly available to encourage research in culturally inclusive AI. The resources are available at: <https://drive.google.com/drive/folders/1UEkhrcaA3-aPQTjTWsIgBC8EbqSS8ntDn?usp=sharing>.

2 Related Works

2.1 Culture in LLMs

Recent studies have examined the sociocultural reasoning of large language models (LLMs), focusing on their adaptability to diverse cultural values and contexts. Researchers have used tools like the World Values Survey and Hofstede’s dimensions to evaluate LLMs’ understanding of social norms and human values (Johnson et al., 2022; Atari et al., 2023; Masoud et al., 2023). Other works have explored cultural artifacts, such as food and art, highlighting LLMs’ knowledge but limited adaptability to user-specific scenarios (Seth et al., 2024b; Li et al., 2024). Efforts to improve adaptability include enabling synthetic personas, but studies reveal gaps in understanding non-Western cultures and persistent biases (Alkhamissi et al., 2024; Durmus et al., 2023). Fine-tuning approaches have been applied to instill social norms and enhance performance on culture-specific tasks like hate speech detection, though probing in regional languages often underperforms compared to monolingual English testing (Dwivedi and Patel, 2024; Shen et al., 2017).

2.2 NLP for Indian Languages

Recent advancements in large language models (LLMs) tailored for India focus on enhancing multilingual capabilities and addressing the country’s linguistic diversity. Notable contributions include IndicGenBench, a multilingual benchmark for 29 Indic languages to evaluate generation capabilities across scripts and language families, promoting research on underrepresented languages (Singh et al., 2024). MuRIL (short for Multilingual Representations for Indian Languages), a language model

specifically designed for Indian languages, effectively handles transliterated and code-mixed text commonly found in informal settings (Khanuja et al., 2021). The LoFTI benchmark (acronym for Localization and Factuality Transfer to Indian Locales) evaluates LLMs’ ability to provide localized and factual responses in Indian contexts (Simon et al., 2024). Additionally, studies on translation capabilities of LLMs between English and 22 Indian languages highlight advancements in multilingual tasks through in-context learning and fine-tuning (Patel et al., 2023).

2.3 Collaborative Dataset Creation

Participatory research involves those impacted by technology in its design and evaluation. While widely used in HCI, such methods are underutilized in NLP (Diddee et al., 2022). Datasets like DOSA (short for “Dataset of Social Artifacts”) Seth et al. (2024a) and CVQA (acronym for “Culturally-diverse Multilingual Visual Question Answering”) Romero et al. (2024) highlight the importance of collaboration, using local input to capture cultural artifacts and diverse insights. Games with a purpose (GWAP) have proven effective for eliciting implicit knowledge (Balayn et al., 2022; Von Ahn et al., 2006). Building on this, we use surveys and a flexible adaptation of the Taboo game to collect cultural artifacts and their significance (Stephenson, 2023), enabling richer and more context-specific data.

3 Development of SANSKRITI

3.1 Data Collection

Data Sources: The creation of the *SANSKRITI* benchmark was a meticulous, multi-phase endeavour designed to ensure comprehensive coverage and high-quality content. The dataset was constructed by sourcing data from six diverse and reliable platforms: Wikipedia², Ritirivaz³, Holidify⁴, Artsandculture⁵, and Timesofindia⁶. These sources were carefully selected to capture the rich tapestry of India’s cultural heritage with authenticity and depth.

Wikipedia: A well-documented and detailed resource, Wikipedia served as the cornerstone for

verified knowledge across diverse domains such as history, arts, languages, and traditions (see appendix for the full domain list).

Ritirivaz: Ritirivaz contributed region-specific insights into Indian customs, rituals, and traditions, providing a valuable perspective on unique practices.

Holidify: Focused on travel, Holidify provided information on festivals, cuisines, and landmarks, effectively linking cultural elements to their geographical origins.

Arts and Culture: Arts and culture enhanced the dataset with visually engaging explorations of art, music, and heritage, complementing textual data with multimedia depth.

Times of India: As a leading news outlet, Times of India provided a contemporary angle, covering cultural events, trends, and regional highlights.

Data Organization: These sources collectively provided a robust repository of historical, traditional, and modern cultural insights. The information was systematically organized into a structured format (“state name”:“attribute”:“scrapped data related to the attribute and state”) to facilitate efficient processing and analysis, ensuring a balanced and authentic representation of India’s cultural diversity.

3.2 Annotation Process

Questions Formation: After collecting the raw data, we curated a structured list of cultural categories—such as arts, festivals, cuisines, music and dance, languages, histories, and traditions (full list in the appendix 7.1)—inspired by Seth et al. (2024b); Romero et al. (2024). These categories were chosen to capture the multifaceted essence of India’s cultural heritage, ensuring comprehensive coverage and balanced representation. The questions in this benchmark are designed as **fact-based** multiple-choice queries.

Cultural Dimensions: By organizing the dataset around these key dimensions, we enable detailed analyses and facilitate applications that require a deep understanding of India’s diverse cultural landscape, significantly enhancing the dataset’s utility.

Team Structure: To generate a comprehensive and high-quality question dataset, we employed a team of forty members, divided into four specialized sub-teams of ten members each. Each sub-team focused on one of the following question

²<https://www.wikipedia.org>

³<https://www.ritirivaz.com>

⁴<https://www.holidify.com>

⁵<https://artsandculture.google.com>

⁶<https://timesofindia.indiatimes.com>

types:

- Association Prediction:** Identify the culture referenced in a statement.

- Country Prediction:** Determine the country based on the statement.

- General Awareness Prediction:** Formulate factual questions.

- State Prediction:** Identify the specific Indian state referenced in a statement.

Quality Assurance: This focused division of labour ensured both efficiency and thoroughness. A total of 21,853 questions were created across these categories, with each sub-team dedicated to its assigned type to maintain quality and minimize redundancy.

Guidelines: To ensure consistency and quality, clear and detailed guidelines were established, outlining definitions for each question type, attributes, sample questions, and cultural considerations. Teams were instructed to prioritize accuracy, diversity, and relevance while avoiding ambiguity or excessive complexity. Outputs from each sub-team underwent cross-validation by another sub-team, fostering collaboration to resolve ambiguities or inaccuracies. A final review ensured adherence to all guidelines.

Team Training and Distribution: Team members, selected for their expertise in linguistics, Indian culture, or related fields, underwent a brief training session on the dataset’s objectives, question types, and cultural sensitivity to maintain a unified approach. Questions were evenly distributed across four types, with minor variations emerging naturally during creation. Redundancy was mitigated by cross-checking for similar content, ensuring diverse coverage of cultural elements.

Cultural Sensitivity and Ethical Considerations: To promote equitable representation, cultural references spanned all Indian states and union territories, with a particular focus on less-represented regions and unique artifacts. Ethical guidelines emphasized cultural sensitivity, avoided stereotypes, and encouraged respectful and inclusive question framing. The cross-validation process further identified and addressed potential biases or misrepresentations, ensuring a balanced and culturally respectful dataset.

Compensation Structure: Lastly, each annotator was compensated for their contributions. The payment structure was designed to reflect the effort involved in each task: annotators were paid 1.20

USD for creating every 10 questions and 0.60 USD for verifying every 10 questions.

<p>TEXT SOURCE: Chhattisgarh is an amalgamation of various tribes and each has its own dance forms. Most of the folk dances of Chhattisgarh are performed as a part of rituals in reverence to Gods and changing seasons. Each dance forms have special costumes and accessories and is perfect synchronization which will leave you mesmerized. Some of the popular dance forms are Salla Dance which is a stick dance and performed during harvest season. It is famous among the people of Sarguja, Chhindwara, and Baitul districts. Karma is a folk dance performed during the autumnal festival and is performed by tribal groups like Gonds, Baigas, and Oraons. Sua Nacha or Suwa Folk dance is a parrot dance performed during the occasion of Goura Marriage, a month in advance of the festival Diwali. Raut Nacha is celebrated after Diwali and the cowherds of Chhattisgarh are the main artist of the dance. Panthi dance is performed by the Satnami community of Chhattisgarh.</p>	<p>State: Chhattisgarh Attribute: Dance and Music</p>
	<p>Ques: Which dance form is traditionally associated with harvest celebrations in Chhattisgarh? Options: Bhangra; Salla Dance; Kathak; Bharatanatyam</p>
	<p>Ques: Chhattisgarh's folk dances are often performed as a part of rituals honoring? Options: Historical figures; Natural elements; Guardians of the village; Gods and changing seasons</p>
	<p>Ques: The occasion of Goura Marriage is celebrated with which of the following dances in Chhattisgarh? Options: Bihu; Sua Nacha; Hula; Lathi Dance;</p>
	<p>Ques: The Salla Dance in Chhattisgarh is characterized by the use of which of the following? Options: Hand gestures; Sticks; Cloth fans; Masks</p>
<p>Ques: The cowherds in Chhattisgarh perform which of the following dances? Options: Bhangra; Panthi; Raut Nacha; Garba</p>	

Figure 1: Illustration showcasing crafting questions on Chhattisgarh’s dance and music Heritage.

Excerpt of Process: Figure 1 illustrates an excerpt of the process used to create the questions shown on the right-hand side of the image, derived from the source text on the left. This process involved identifying key details, themes, and cultural elements described in the passage about Chhattisgarh’s dance forms. Each question was carefully crafted to test specific factual knowledge from the text. For example, the association of the Salla Dance with the harvest season, the ritualistic purpose of folk dances, and the connection of Sua Nacha to Goura Marriage were directly drawn from explicit references in the source. Plausible distractors were included for each question, while the correct answer, highlighted in red, closely adhered to the source text. The question formulation prioritized clarity and relevance, aligning each item with distinct phrases or concepts (e.g., “stick dance,” “rituals honouring gods and seasons,” and “cowherds performing Raut Nacha”).

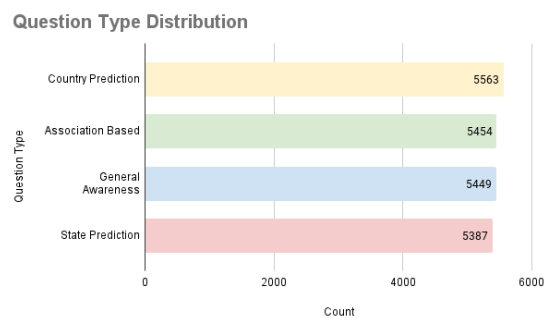


Figure 2: Distribution of Question Types Showing Balanced Representation Across Categories

3.3 Statistical details of SANSKRITI

Our multidimensional dataset is analyzed through three key visualizations: a bar chart of question types (Figure 2), a state-wise India map (Figure 3), and a bar chart of attribute-wise question counts (Figure 4), offering a concise overview of its structure, diversity, and coverage.

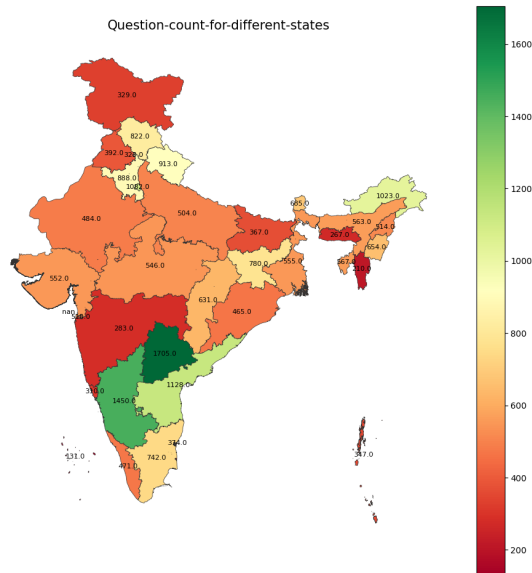


Figure 3: A map of India illustrating the distribution of questions across all states and union territories.

Distribution of Questions across States and Union Territories: As shown in Figure 3, the map visualizes the distribution of question counts across different Indian states using a colour gradient where green represents higher counts and red represents lower counts. States like Telangana, Karnataka, and Andhra Pradesh have relatively higher question counts, signifying high engagement or activity. Conversely, Jammu and Kashmir, Maharashtra, and smaller northeastern states like Tripura show lower counts, indicated by red and orange shades. The range spans 300 to over 800, with central and southern states exhibiting stronger performance, while northern and northeastern regions reflect mixed trends. We must highlight that we have treated Dadra, Nagar Haveli, Daman, and Diu as a single entity in our work. This decision was made because, when considered separately, each region had a relatively small number of questions, whereas combining them resulted in a more significant count.

Distribution of Questions across the Question Categories: The bar chart shown in Figure 2 illustrates the distribution of various question

types by their counts. Among the four categories, "Country Prediction" has the highest representation, closely followed by "Association Based" and "General Awareness," which are almost equal in count. "State Prediction" has slightly fewer instances but remains comparable to the other categories. Overall, the chart shows a balanced distribution of question types, with no significant disparity among them, indicating a well-rounded approach to creating diverse question types in the dataset.

Distribution of Questions across the Attributes: The bar chart presented in Figure 4, shows the distribution of attributes based on their count. It reveals that "Tourism", "Festivals," and "History" are the most frequently represented attributes, indicating a strong focus on cultural and historical themes. Other significant attributes include "Art," "Cuisine," and "Dance and Music," showcasing an emphasis on elements reflecting regional heritage and traditions. "Costume" and "Cultural Common Sense" also have notable representation, highlighting their importance in the context of cultural identity. Attributes like "Languages" and "Rituals" are moderately represented, while "Religion" has relatively lower prominence. The least represented attributes include "Medicine," "Transport," "Sports," and "Nightlife," suggesting these are less explored or emphasized within the dataset. Overall, the chart prioritizes cultural and historical aspects over contemporary or niche topics.

Word Cloud Representation: Last but not least, Figure 5 presents word cloud visualizations for the four question types: Association Based, General Awareness, Country Prediction, and State Prediction. The Association Based category highlights key cultural and linguistic terms, while General Awareness emphasizes geographical knowledge through prominent state names. Country Prediction reflects national identity and traditions, whereas State Prediction focuses on regional and cultural themes. These word clouds collectively underscore the dataset's strong emphasis on cultural, religious, and geographical associations.

4 Experimental Section

4.1 Models

To conduct a comprehensive evaluation of our benchmark, *SANSKRITI*, we performed an extensive assessment across a diverse range of language models. Our evaluation included prominent Large Language Models (LLMs), such as Mistral-

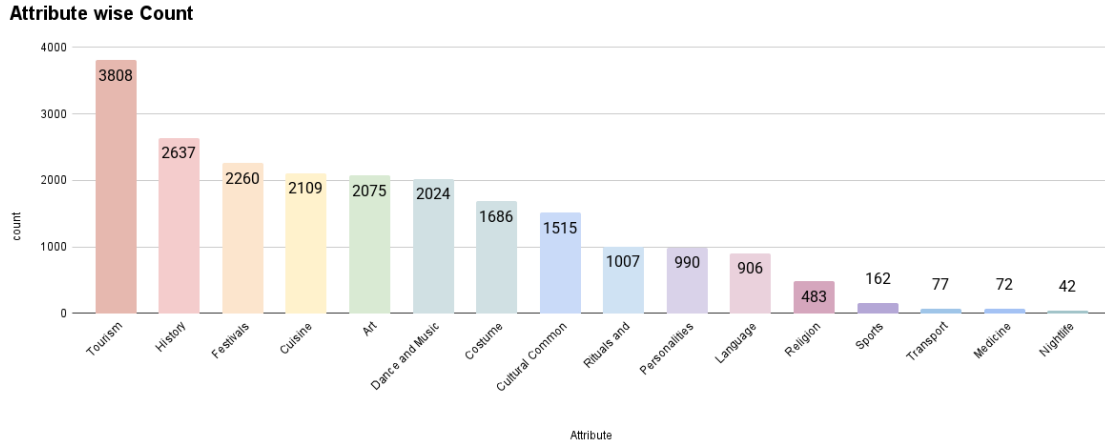


Figure 4: Distribution of Question Types across various Attributes.

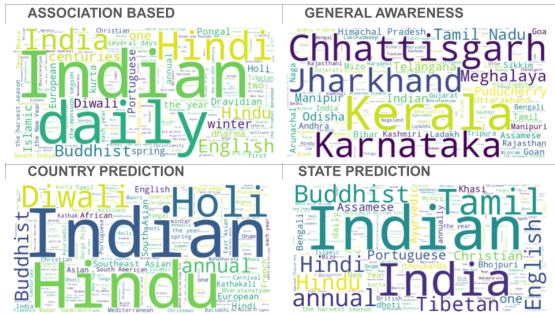


Figure 5: Word Cloud Representation of Question Categories

7B-Instruct, LLAMA-3.1-70B-Instruct, Qwen-2.5-72B-Instruct, and Phi3-medium-4k-Instruct. Additionally, we tested several Small Language Models (SLMs), including Gemma-2B-Instruct, Qwen2-1.5B-Instruct, LLAMA-3.2-3B-Instruct, and SmolLM-1.7B-Instruct. For Indic LLMs, we evaluated Navrasa-2.0 and OpenHathi-Instruct. Furthermore, we incorporated the proprietary model GPT-4o into our evaluation for a holistic comparison.⁷

4.2 Evaluation Setup

We conducted a zero-shot evaluation of cultural aspect-related questions, categorizing them into four main groups: **1. Country Prediction** : Prediction of the country based on cultural aspects, **2. State Prediction**: Prediction of the state in India is based on cultural aspects. **3. General Knowledge**: General awareness questions based on cultural attributes. **4. Association-based Questions**

⁷In the work GPT and GPT-4o are used interchangeably. Both mean the same in this context.

All questions followed a **Multiple-Choice Question (MCQ)** format, with options (A, B, C, D). The MCQs used to evaluate the cultural knowledge of the models are as follows: **1. Country Prediction**: (e.g., “Location: Unknown. Question: Which country is associated with cultural_aspect? Options: options Short Answer:”). **2. State Prediction** : (e.g., “Location: India. Question: Which Indian state is known for cultural_aspect? Options: options Short Answer:”). **3. General Knowledge** : (e.g., “Question: general_cultural_question? Options: options Short Answer:”). **4. Association-Based Questions**: (e.g., “Question: The cultural_entity is most closely associated with which cultural_context? Options: options Short Answer:”). For evaluation, we used **accuracy** as the sole metric. During inference, open-source models were run using 16-bit floating-point precision with **greedy decoding**, while proprietary models were accessed via their respective APIs. The models produced output probabilities for each option, and the option with the highest probability was selected as the prediction. This approach ensured consistency across all evaluations.

5 Discussion on Results

5.1 Main Results

The overall performance on the *SANSKRITI* dataset across various LLMs, SLMs, and ILMs is presented in Table 1.

Performance of LLMs : GPT-4o demonstrated the best performance on the *SANSKRITI* dataset. Among the open-source LLMs, **LLAMA-3.1B-70B-Instruct** achieved the highest perfor-

QuestionType	Gemma-2-2b	Qwen2-1.5B	Llama-3.2-3B	SmolLM-1.7B	Navarasa-2.0	OpenHathi-7B	Phi-3-medium-4k	Mistral-7B	Llama-3.1-70B	Qwen2.5-72B	GPT-4o
Association Prediction	0.53	0.75	0.43	0.18	0.26	0.3	0.71	0.58	0.85	0.80	0.82
State Prediction	0.40	0.51	0.50	0.18	0.28	0.2	0.7	0.61	0.78	0.72	0.80
GK Prediction	0.20	0.82	0.79	0.16	0.56	0.35	0.85	0.81	0.93	0.94	0.96
Country Prediction	0.81	0.9	0.36	0.13	0.5	0.43	0.84	0.8	0.88	0.92	0.93

Table 1: Performance comparison of various LLMs and SLMs across tasks including Association Prediction, State Prediction, GK Prediction, and Country Prediction. The table presents the accuracy achieved by each model.

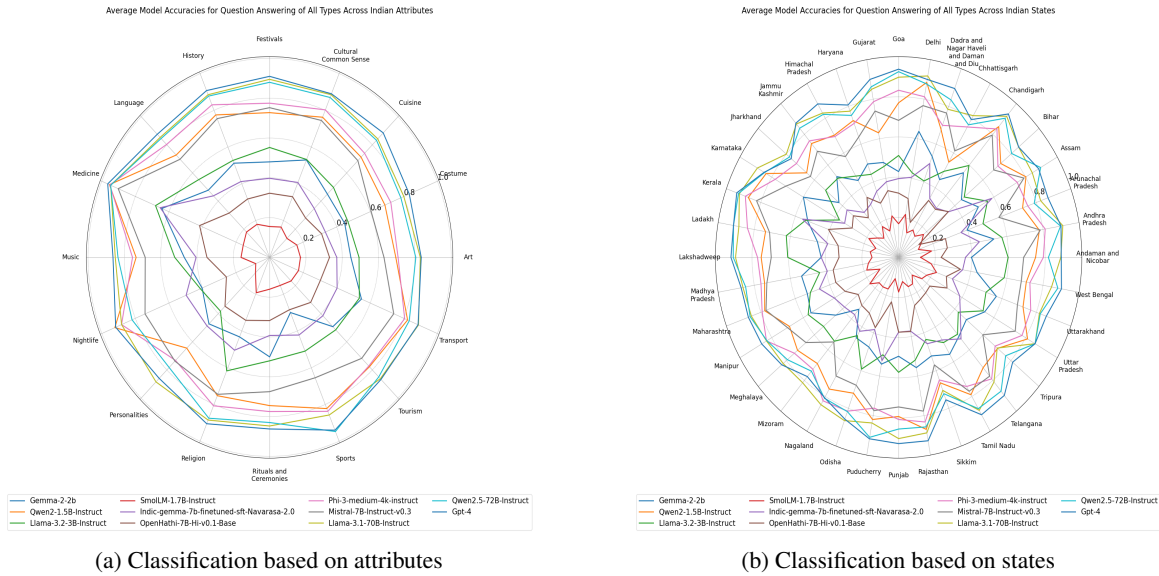


Figure 6: Average accuracy of language models on the *SANSKRITI* dataset classified on the basis of attributes and states

Model	Attribute Score
Gemma-2-2b	0.48
Qwen2-1.5B-Instruct	0.74
Llama-3.2-3B-Instruct	0.52
SmolLM-1.7B-Instruct	0.16
Navarasa-2.0	0.40
OpenHathi-7B-Instruct	0.32
Phi-3-medium-4k-Instruct	0.77
Mistral-7B-Instruct-v0.3	0.70
Llama-3.1-70B-Instruct	0.86
Qwen2.5-72B-Instruct	0.84
GPT-4o	0.87

Table 2: Average accuracy of various language models on the entire *SANSKRITI* dataset.

mance with a score of 0.86, followed by **Phi-3-medium-4k-Instruct** and **Mistral-7B-Instruct**, both showed comparable performance scores of 0.77 and 0.7, respectively.

Performance of SLMs: Among SLMs, **Qwen2-1.5B-Instruct** emerged as the best-performing SLM. Interestingly, it is also almost comparable to a few notable LLMs in our evaluation, highlighting that some SLMs possess better domain-specific knowledge than even larger LLMs. The performances of **Llama-3.2-3B-Instruct** and **SmolLM-1.7B-Instruct** were comparable, while **Gemma-**

2-2B with a score of 0.48 delivered the weakest performance among the SLMs.

Performance of ILM: Among the ILMs evaluated, **OpenHathi** emerged as the best performer with a score of 0.32. Surprisingly, **Navrasa-2.0**, despite being an ILM delivered the weakest performance across all the language models evaluated on our benchmark *SANSKRITI* dataset.

5.2 Performance across Cultural Attributes

To assess the fine-grained capabilities of LLMs, SLMs, and ILMs in addressing questions related to various attributes of Indian culture, we present their accuracy comparison across different attributes in Figure-15b a. Our analysis reveals that for certain attributes, such as religion, medicine, and cultural common sense, the performance of all language models was notably high. However, for attributes like costumes, cuisines, and art, the language models generally struggled. Additionally, some trends were specific to certain LLMs—for instance, the performance of **Gemma-2-2b** dropped significantly for sports compared to other LLMs.

5.3 Performance across Different States

Similarly, we conducted a fine-grained analysis of language models across questions from various states and union territories of India. A notable trend was that language models struggled with questions pertaining to North-Eastern states such as Sikkim, Arunachal Pradesh, and Tripura. Poor performance was also observed for states like Bihar and Jharkhand. Conversely, language models performed relatively well for states like Delhi and Maharashtra. An interesting observation was that states with globally recognized cities tended to yield better results compared to others.

5.4 Performance across Question Types

We also evaluated the performance across different question types in *SANSKRITI*. The evaluation revealed several noteworthy trends. Across all language models, the highest accuracy was achieved on **GK**-based questions, whereas **State Prediction** questions recorded the lowest accuracy. Overall, **GPT-4.0** emerged as the top-performing model in our evaluation. However, for **Association-Based** questions, **LLAMA-3.1B-70B** outperformed all other models, showcasing its strengths in this specific category. Remarkably, **Qwen2-1.5B-Instruct**, despite its smaller size, performed almost at par with the state-of-the-art **GPT-4.0** on the **Country Prediction** task, highlighting its efficiency and potential.

5.5 Error Analysis

To evaluate the strengths and limitations of the best-performing model, GPT-4o, we present an error analysis by examining questions grouped into correctly and incorrectly answered sets. This analysis underscores the model’s proficiency in leveraging strong semantic associations while shedding light on challenges such as limited contextual understanding, ambiguous options, and gaps in the representation of subtle knowledge within the training data.

As illustrated in Figure 7, the LHS showcases the model’s correct predictions, which likely arise from the strong semantic associations between keywords (e.g., “Pandal”, “Yoga”) and their contexts—areas well-represented in the training data. These questions are characterized by culturally specific, unambiguous topics with clear, non-overlapping options, enabling the model to effectively identify distinct patterns and contextual cues. Additionally, the clar-

<p>Ques: What does the term "Pandal" refer to in the context of Durga Puja? Options: A type of food; Temporary structure for worship; Traditional attire; A sacred text Predicted: Temporary structure for worship Correct: Temporary structure for worship</p>	<p>Ques: What is the primary source of livelihood for many local communities in Ladakh, affected by tourism? Options: Agriculture; Fishing; Craftsmanship and hospitality; Mining Predicted: Agriculture Correct: Craftsmanship and hospitality</p>
<p>Ques: A long, draped cloth reaching the knees known as Rignai is a staple in the traditional wear of this state. Options: Tripura; Tamil Nadu; Bihar; Nagaland Predicted: Tripura Correct: Tripura</p>	<p>Ques: Which city is often regarded as the cultural capital of Punjab? Options: Amritsar; Chandigarh; Delhi; Lahore Predicted: Lahore Correct: Amritsar</p>
<p>Ques: The tradition of yoga, believed to bring inner peace and holistic health, has its roots in this country. Options: China; India; Japan; Thailand Predicted: India Correct: India</p>	<p>Ques: Qawwali, a musical form associated with Sufi traditions, is particularly popular in which region? Options: Haryana; Rajasthan; Uttar Pradesh; Gujarat Predicted: Rajasthan Correct: Uttar Pradesh</p>
<p>Ques: Participating in a pilgrimage to a sacred river is considered a vital spiritual practice. Options: Ancient Egyptian Culture; Indian Hindu Culture; Celtic Culture; Indigenous Australian Culture Predicted: Indian Hindu Culture Correct: Indian Hindu Culture</p>	<p>Ques: The term "Himachal Pradesh" literally refers to what in Sanskrit? Options: Land of the Himalayas; Gateway to the North; Land of Mountains; Abode of the Gods Predicted: Abode of the Gods Correct: Land of the Himalayas</p>

Figure 7: The LHS displays the correctly answered questions, while the RHS highlights the incorrectly answered ones.

ity and relevance of the phrasing further aid in accurate classification, minimizing confusion.

Similarly, the RHS highlights the model’s incorrect predictions, which can be attributed to several factors. For instance, questions involving nuanced cultural or regional knowledge (e.g., primary livelihood in Ladakh, cultural capital of Punjab) require deeper contextual understanding, which the model may lack due to limited or biased training data. Ambiguity in certain options (e.g., “Land of the Himalayas” vs. “Abode of the Gods”) further compounds the challenge, as similar terms can confuse the model. Additionally, less explicit semantic associations between keywords and correct answers, coupled with overlapping or less distinctive options, reduce the model’s ability to identify the correct patterns effectively.

6 Conclusion

In this work, we introduced *SANSKRITI*, a benchmark to evaluate language models’ understanding of India’s cultural diversity through over 21,000 curated question-answer pairs spanning 28 states and 8 union territories. *SANSKRITI* meets the need for culturally nuanced datasets by covering attributes like arts, festivals, and cuisines. Model evaluations revealed significant gaps, especially for certain states, highlighting biases likely rooted in training data limitations. *SANSKRITI* dataset ensured quality, precision and cultural sensitivity, setting a new standard for inclusive AI research. By releasing *SANSKRITI* publicly, we aim to advance the development of AI systems that are culturally aware and capable of serving diverse populations effectively. Looking forward, we plan to expand *SANSKRITI* by adding more attributes to capture nuanced aspects of Indian culture, such as regional folklore and ecological heritage. Additionally, we

aim to build a visual question-answering dataset and a multilingual version of **SANSKRITI**, enabling broader accessibility and applications across diverse linguistic and cultural contexts.

Limitations

Although our study offers the most comprehensive evaluation of language models for Indian cultural knowledge, it has several notable limitations:

1) Limited Scope of Cultural Attributes: This study covers only sixteen cultural attributes, which may not fully represent Indian culture. The question types also may not capture its full complexity. Future work will expand attributes and incorporate diverse QA formats, including True/False and adversarial questions.

2) Lack of State-Specific Multilingual Queries: While our queries are multicultural, they do not currently include state-specific multilingual questions. Incorporating such queries is a key priority for the immediate extension of this work, as it will provide a more robust evaluation of language models across diverse linguistic contexts.

3) Absence of Visual Question-Answering (VQA) Tasks: The dataset does not currently address visual question-answering tasks, which have become increasingly important with recent advancements in visual language models (VLMs). Future work will seek to incorporate VQA to further enhance the dataset’s applicability and relevance.

4) Limited Contextual Clarity: Although we aimed to include attributes that are uniquely representative of Indian culture, certain examples in the benchmark may involve cultural elements that are somewhat ambiguous, as similar terms or practices can also be associated with other countries. In such cases, where the context is not detailed enough to clearly isolate India, we carefully design the multiple-choice questions so that the distractor options do not reference the attribute mentioned in the prompt. We then replace the original correct answer with the option that most closely aligns with it among the four choices.

5) Lack of Diversity in Questions: The questions in this benchmark are multiple-choice, fact-based questions that do not explicitly require reasoning or causal understanding to answer them.

Ethics Statement

Data Collection and Bias Mitigation: The data used in the development of **SANSKRITI** was sourced from publicly accessible platforms, as detailed in Section 3.1. While these platforms were carefully chosen to ensure authenticity and diversity, there remains a possibility of underrepresentation of certain cultural domains or regions. Despite its limitations, the **SANSKRITI** dataset marks a significant milestone in establishing a standardized and inclusive benchmark for evaluating India’s cultural knowledge. By encompassing all 28 states and 8 union territories, and featuring an extensive collection of 21,853 questions, it stands as the largest dataset dedicated to testing fine-grained cultural understanding of India.

Human Annotation and Cultural Sensitivity: Human annotators played a key role in crafting and validating the dataset to ensure high quality and cultural accuracy. A diverse group of 40 annotators—selected for their expertise in Indian culture, linguistics, or related fields—included native and bilingual speakers from various Indian states. About 75% were native speakers of at least one Indian language, and 80% had lived over 15 years in regions where their primary language is spoken, ensuring strong cultural grounding. Annotators, aged 20 to 45, underwent training on dataset objectives, question categories, and cultural sensitivity. The annotation process was collaborative, with outputs cross-validated by a separate team to ensure consistency and mitigate bias. Ethical guidelines emphasized inclusivity and the avoidance of stereotypes, with any misrepresentations addressed during cross-validation to ensure respectful representation of India’s diverse culture.

Acknowledgments

Raghvendra Kumar gratefully acknowledges support from the Prime Minister’s Research Fellowship (PMRF). Akash Ghosh expresses gratitude to the SERB POWER scheme (SPG/2021/003801), Department of Science and Technology, Government of India. Sriparna Saha acknowledges funding from the Technology Innovation Hub, Vishlesan I-Hub Foundation, IIT Patna (Project No. TIH/CSE/ASMO/05).

References

- Badr AlKhamissi, Muhammad ElNokrashy, Mai AlKhamissi, and Mona Diab. 2024. Investigating cultural alignment of large language models. *arXiv preprint arXiv:2402.13231*.
- Mikel Artetxe et al. 2020. Xquad: A cross-lingual question answering dataset. In *EMNLP 2020*.
- Mohammad Atari, Jonathan Haidt, Jesse Graham, Sena Koleva, Sean T Stevens, and Morteza Dehghani. 2023. Morality beyond the weird: How the nomological network of morality varies across cultures. *Journal of Personality and Social Psychology*.
- Agathe Balayn, Gaole He, Andrea Hu, Jie Yang, and Ujwal Gadiraju. 2022. Ready player one! eliciting diverse knowledge using a configurable game. In *Proceedings of the ACM Web Conference 2022*, pages 1709–1719.
- Emily M. Bender, Timnit Gebru, et al. 2021. On the dangers of stochastic parrots: Can language models be too big? *FAccT 2021*, pages 610–623.
- Su Lin Blodgett, Solon Barocas, et al. 2020. Language (technology) is power: A critical survey of “bias” in nlp. In *ACL 2020*.
- Tom Brown, Benjamin Mann, Nick Ryder, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL 2019*, pages 4171–4186.
- Harshita Diddee, Kalika Bali, Monojit Choudhury, and Namrata Mukhija. 2022. The six conundrums of building and deploying language technologies for social good. In *Proceedings of the 5th ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies*, pages 12–19.
- Esin Durmus, Karina Nyugen, Thomas I Liao, Nicholas Schiefer, Amanda Askill, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, et al. 2023. Towards measuring the representation of subjective global opinions in language models. *arXiv preprint arXiv:2306.16388*.
- Sanjay Kumar Dwivedi and Rahul Patel. 2024. Exploring the intersections: Anthropological insights into studying language and culture. *State Institute of Education, Allahabad*, 30:171–182.
- Akash Ghosh, Arkadeep Acharya, Prince Jha, Sriparna Saha, Aniket Gaudgaol, Rajdeep Majumdar, Aman Chadha, Raghav Jain, Setu Sinha, and Shivani Agarwal. 2024a. Medsumm: A multimodal approach to summarizing code-mixed hindi-english clinical queries. In *European Conference on Information Retrieval*, pages 106–120. Springer.
- Akash Ghosh, Arkadeep Acharya, Sriparna Saha, Gau-rav Pandey, Dinesh Raghu, and Setu Sinha. 2024b. Healthalignsumm: Utilizing alignment for multimodal summarization of code-mixed healthcare dialogues. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 11546–11560.
- Akash Ghosh, Debayan Datta, Sriparna Saha, and Chirag Agarwal. 2025. The multilingual mind: A survey of multilingual reasoning in language models. *arXiv preprint arXiv:2502.09457*.
- Akash Ghosh, Mohit Tomar, Abhisek Tiwari, Sriparna Saha, Jatin Salve, and Setu Sinha. 2024c. From sights to insights: Towards summarization of multimodal clinical documents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13117–13129.
- Rebecca L Johnson, Giada Pistilli, Natalia Menéndez-González, Leslye Denisse Dias Duran, Enrico Panai, Julija Kalpokiene, and Donald Jay Bertulfo. 2022. The ghost in the machine has an american accent: value conflict in gpt-3. *arXiv preprint arXiv:2203.07785*.
- Pranav Joshi et al. 2020. Tydi qa: A benchmark for information-seeking question answering in typologically diverse languages. In *ACL 2020*.
- Simran Khanuja, Diksha Bansal, Sarvesh Mehtani, Savya Khosla, Atreyee Dey, Balaji Gopalan, Dilip Kumar Margam, Pooja Aggarwal, Rajiv Teja Nagipogu, Shachi Dave, Shruti Gupta, Subhash Chandra Bose Gali, Vish Subramanian, and Partha Talukdar. 2021. [MuriL: Multilingual representations for indian languages](#).
- Huihan Li, Liwei Jiang, Jena D Hwang, Hyunwoo Kim, Sebastin Santy, Taylor Sorensen, Bill Yuchen Lin, Nouha Dziri, Xiang Ren, and Yejin Choi. 2024. Culture-gen: Revealing global cultural perception in language models through natural language prompting. *arXiv preprint arXiv:2404.10199*.
- Percy Liang et al. 2022. Towards holistic evaluation of language models. *NeurIPS 2022*.
- Bill Yuchen Lin et al. 2021. Understanding cultural and gender biases in multilingual language models. In *NAACL 2021*.
- Reem I Masoud, Ziquan Liu, Martin Ferianc, Philip Treleaven, and Miguel Rodrigues. 2023. Cultural alignment in large language models: An explanatory analysis based on hofstede’s cultural dimensions. *arXiv preprint arXiv:2309.12342*.
- Vishnu Narayanan. 2020. Festivals of india: Cultural and religious significance. *Journal of Indian Traditions*.
- Long Ouyang, Jeff Wu, Xu Jiang, et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.

Ramesh Patel et al. 2023. Evaluating cultural understanding in multilingual models. *arXiv preprint arXiv:2302.12345*.

David Romero, Chenyang Lyu, Haryo Akbarianto Wibowo, Teresa Lynn, Injy Hamed, Aditya Nanda Kishore, Aishik Mandal, Alina Dragonetti, Artem Abzaliev, Atnafu Lambebo Tonja, et al. 2024. Cvqa: Culturally-diverse multilingual visual question answering benchmark. *arXiv preprint arXiv:2406.05967*.

Agrima Seth, Sanchit Ahuja, Kalika Bali, and Sunayana Sitaram. 2024a. Dosa: A dataset of social artifacts from different indian geographical subcultures. *arXiv preprint arXiv:2403.14651*.

Agrima Seth, Sanchit Ahuja, Kalika Bali, and Sunayana Sitaram. 2024b. **DOSA: A dataset of social artifacts from different Indian geographical subcultures**. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 5323–5337, Torino, Italia. ELRA and ICCL.

Meera Sharma. 2019. *Indian Culture and Heritage: A Comprehensive Guide*. Publisher XYZ.

Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. *Advances in neural information processing systems*, 30.

Sona Elza Simon, Soumen Kumar Mondal, Abhishek Singhania, Sayambhu Sen, and Preethi Jyothi. 2024. Lofti: Localization and factuality transfer to indian locales. *arXiv preprint arXiv:2407.11833*.

Harman Singh, Nitish Gupta, Shikhar Bharadwaj, Dinesh Tewari, and Partha Talukdar. 2024. Indicgenbench: A multilingual benchmark to evaluate generation capabilities of llms on indic languages. *arXiv preprint arXiv:2404.16816*.

Janet Stephenson. 2023. *Culture and Sustainability: exploring stability and transformation with the cultures framework*. Springer Nature.

Ramesh Tripathi. 2019. Cultural diversity of india. *Indian Historical Review*.

Luis Von Ahn, Mihir Kedia, and Manuel Blum. 2006. Verbosity: a game for collecting common-sense facts. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 75–78.

7 Appendix

7.1 Complete List of Attributes and their definitions

- **Rituals and Ceremonies:** Traditions and practices performed during cultural or religious events.

- **Cultural Common Sense:** Shared knowledge and social norms unique to a culture.
- **History:** Chronological account of significant events and developments from the past.
- **Tourism:** Exploration of culturally or historically significant destinations.
- **Cuisine:** Traditional foods, cooking styles, and regional delicacies.
- **Dance and Music:** Artistic expressions through movement and sound representing cultural identity.
- **Costume:** Traditional attire and accessories specific to a region or culture.
- **Language:** Spoken and written forms of communication unique to a culture or region.
- **Art:** Visual and creative works, including paintings, sculptures, and crafts.
- **Festivals:** Celebrations marked by cultural or religious significance.
- **Religion:** Belief systems and practices centered around faith and spirituality.
- **Medicine:** Traditional healing practices and medical knowledge of a culture.
- **Transport:** Modes of travel and transportation characteristic of a region.
- **Sports:** Physical activities and games enjoyed for competition or recreation.
- **Nightlife:** Cultural activities and entertainment available during the evening.
- **Personalities:** Notable individuals who have significantly influenced cultural or historical narratives.

7.2 Question Formulation from the Source text and the Rationale behind it.

7.3 Word Clouds based on questions classified on different cultural attributes and Results on different kind of questions

Source text for attribute “History” and state “Bihar”

Bihar’s history can be traced back to the Prehistoric Period and is very ancient. The earliest history can be traced to the Hindu epic of Ramayana. Mithila was the birthplace of Lord Ram’s wife, Sita. The author of The Ramayan, Maharishi Valmiki lives in ancient Bihar. Mahajanpada period and the rule of the Magadhan empire bring glory to ancient Bihar. Pataliputra (Patna) was the capital of ancient India’s powerful Magadha kingdom. Bihar has experienced various invasions from different dynasties. Two of India’s most glorious dynasties, Mauryas (321 -185 BCE) and Guptas (320 to 550 CE) flourished in the ancient Bihar region which was then known as Magadh. The Mauryan emperor, Ashoka (born c.304 BCE, died c. 232 BCE), who was born in Pataliputra (Patna) is believed to be one of the greatest rulers in the history of the world. The Great King Ashoka of the Mauryan dynasty whose empire spread across South Asia, had his capital in Pataliputra which is present Patna. It was a prosperous center of India’s ancient civilization. Bihar has been associated with Chandragupta Maurya and Emperor Ashoka one of the greatest empires in India. There were 19 thousand Buddhist monasteries in Magadha during Emperor Ashoka. The inscriptions of Ashok, his Dharma, and other features like the Ashoka pillar have been incorporated into the independent Republic of India. The Dharma chakra is incorporated into the national flag of India, The Indian tricolor and the figure of four lions atop a pedestal, with the inscription of a wheel, were adopted as the Official Seal of the independent Republic of India. The great spiritual way of Buddhism originated and thrived in ancient Bihar, as Buddha attained his enlightenment in Bodhgaya. Therefore the region is full of remains of the monasteries known as Stupa. Jain leader Mahavira also belonged to this region and attained his Nirvana at Pawapuri. In the 12th century, the invasion of Muhammad bin Khilji resulted in the Nalanda and Taxila schools destruction along with thousands of Buddhist monks. Guru Govind Singh, the tenth and last Guru of the Sikhs, was born in Patna. The Har Mandir Takht built by Maharaja Ranjit Singh to commemorate his birthplace is regarded as one of the five ‘Akal Takht’s by the Sikhs. In medieval times, Bihar was at its peak during the reign of Sher Shah Suri who founded the city of Patna on the ancient land of Pataliputra. In the reign of the great Afghan ruler Sher Shah Suri ruled from Patna and built the Grand Trunk Road the longest road in India. Bihar saw a span of five years of good governance. During the freedom struggle, Bihar has some greatest revolutionizers like Khudiram Bose, Prafulla Chaki, and Chandrasekhar Azad. During British rule, Mahatma Gandhi started his first experiment of the Satyagraha movement from the Champaran region against the oppression of the indigo farmers by the Britishers. Bihar remained a part of the Bengal Presidency of British India until December 1911, when it was separated from the Bengal Presidency while Bihar and Orissa comprised a single province. The constitution of the united Bihar and Orissa was changed to separate the state of Bihar in 1936. The present form of Bihar came into existence on 01 November 1956. In 2000, Jharkhand was made a separate state from Bihar.

Questions formulated using above source text and the rationale behind it

Question	Option A	Option B	Option C	Option D
Q: This state was home to one of the greatest Mauryan emperors, who is celebrated for spreading Buddhism across Asia.	West Bengal	Bihar	Andhra Pradesh	Tamil Nadu

Q: The site where Buddha attained enlightenment is a significant pilgrimage destination in this state.	Punjab	Assam	Bihar	Himachal Pradesh
Q: This state features numerous ancient Buddhist monasteries known as stupas, signifying its importance in the spread of Buddhism.	Uttarakhand	Bihar	Odisha	West Bengal
Q: The region where the ancient rule of the Guptas was established is known in modern times as this state.	Chhattisgarh	Bihar	Madhya Pradesh	Haryana
Q: The ancient ruins of Nalanda University can be found in this country, recognized as one of the world's first residential universities.	Thailand	India	Pakistan	Nepal
Q: Known for the Maurya and Gupta Empires, this country has a rich historical legacy influencing art and politics.	India	Iran	Egypt	Greece
Q: The ancient city of Pataliputra, now known as Patna, was the capital of a prosperous empire in this country.	India	Pakistan	Afghanistan	Bangladesh
Q: This country has a significant tradition of yoga and meditation, deeply rooted in its historical and cultural practices.	Japan	India	China	Thailand

The rationale for each question formulated based on the source text:

1. **Q: This state was home to one of the greatest Mauryan emperors, who is celebrated for spreading Buddhism across Asia.**
 - *Rationale:* The question highlights Bihar's association with Emperor Ashoka, a prominent figure in spreading Buddhism. Ashoka's legacy is a key historical aspect mentioned in the source text, making Bihar the correct answer.
2. **Q: The site where Buddha attained enlightenment is a significant pilgrimage destination in this state.**

- *Rationale:* The question emphasizes Bodhgaya, a globally recognized Buddhist pilgrimage site in Bihar, where Buddha attained enlightenment. This historical and spiritual connection aligns with Bihar's cultural identity.
3. **Q: This state features numerous ancient Buddhist monasteries known as stupas, signifying its importance in the spread of Buddhism.**
- *Rationale:* The mention of numerous Buddhist monasteries and stupas in Bihar in the source text is the basis for this question. It underlines the state's historical role in the propagation of Buddhism.
4. **Q: The region where the ancient rule of the Guptas was established is known in modern times as this state.**
- *Rationale:* The Gupta Empire, which flourished in ancient Bihar (then Magadha), is central to the state's historical narrative. The question connects ancient historical events to the modern-day state.
5. **Q: The ancient ruins of Nalanda University can be found in this country, recognized as one of the world's first residential universities.**
- *Rationale:* Nalanda University, a globally renowned center of learning in ancient Bihar, is integral to India's cultural and educational heritage. The question ensures the recognition of India's historical contributions.
6. **Q: Known for the Maurya and Gupta Empires, this country has a rich historical legacy influencing art and politics.**
- *Rationale:* This question broadens the scope by linking India's historical legacy to its global impact, specifically focusing on the Maurya and Gupta empires, which are explicitly mentioned in the source text.
7. **Q: The ancient city of Pataliputra, now known as Patna, was the capital of a prosperous empire in this country.**
- *Rationale:* Pataliputra's transformation into modern-day Patna is a significant historical detail. The question connects ancient history with contemporary geography, emphasizing Bihar's importance.
8. **Q: This country has a significant tradition of yoga and meditation, deeply rooted in its historical and cultural practices.**
- *Rationale:* Although not directly tied to Bihar, the question uses India's broader historical narrative of yoga and meditation to highlight its cultural heritage. This aligns with the overarching theme of India's historical and cultural significance.

The annotators ensured the questions are grounded in the text, contextually relevant, and highlight key historical aspects while diversifying the focus to emphasize Bihar's unique contributions and India's broader cultural legacy.

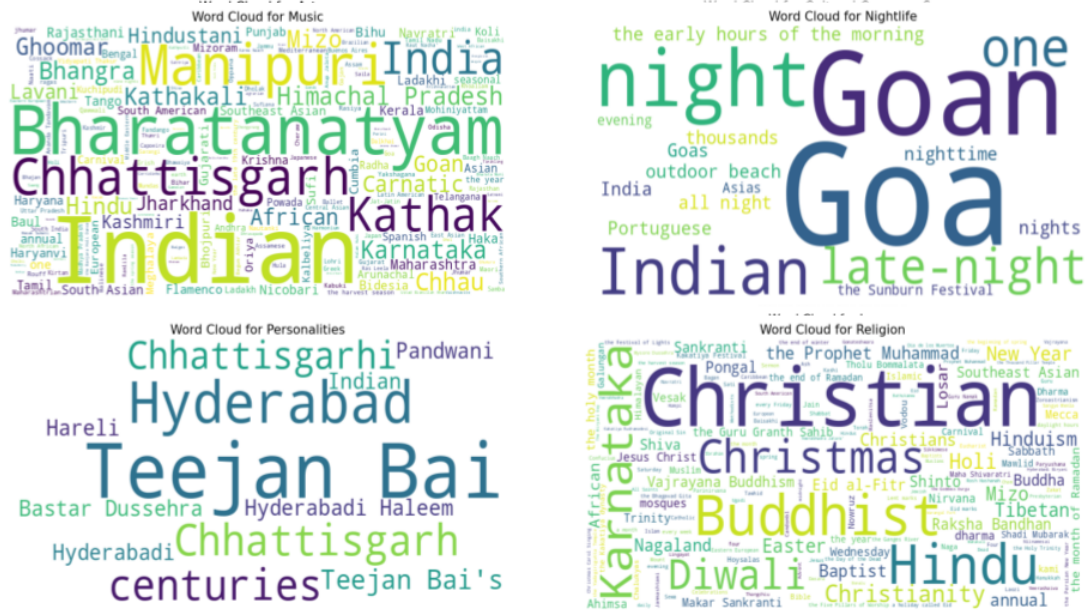


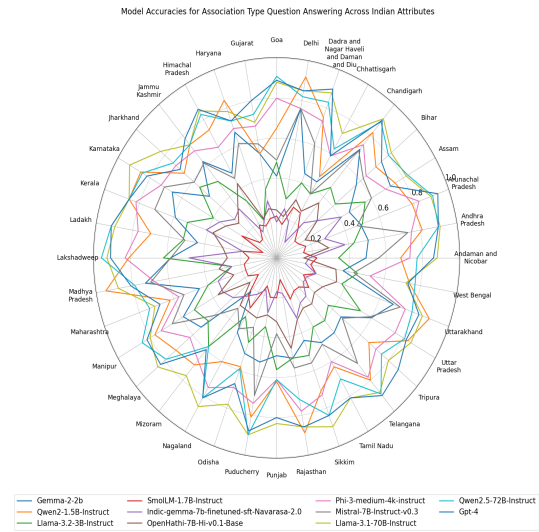
Figure 10: Word Cloud Representation based on attributes Categories like transport, sports, tourism and ceremonies



Figure 11: Word Cloud Representation based on attributes Categories like history ,language ,medicine ,festivals.

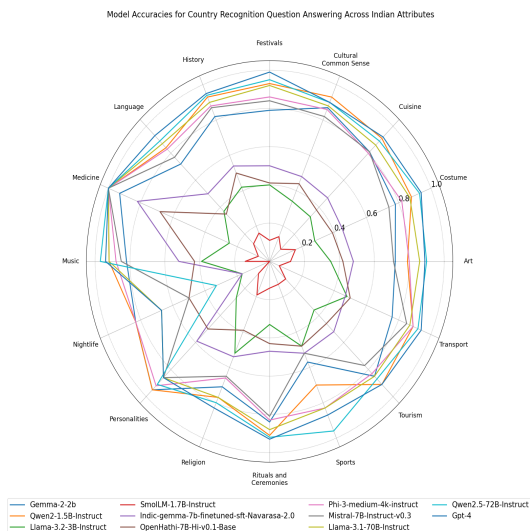


(a) Performance of Language models on Association kind of Questions classified on the basis of attributes

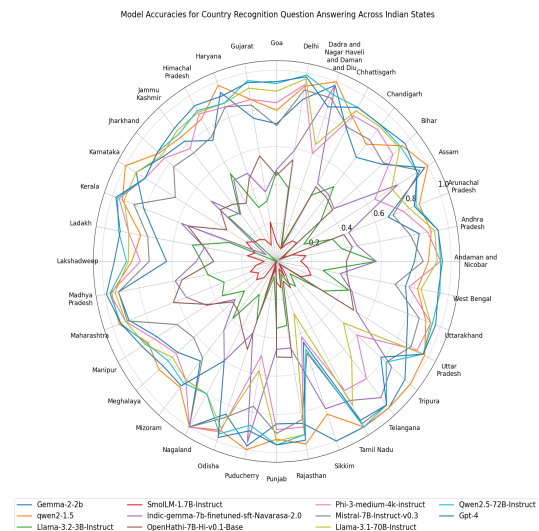


(b) Performance of Language models on Association kind of Questions classified on the basis of states.

Figure 12: Performance of language models on the data points that falls under association kind of questions in *SANSKRITI* that are further classified on the basis of attributes and states.

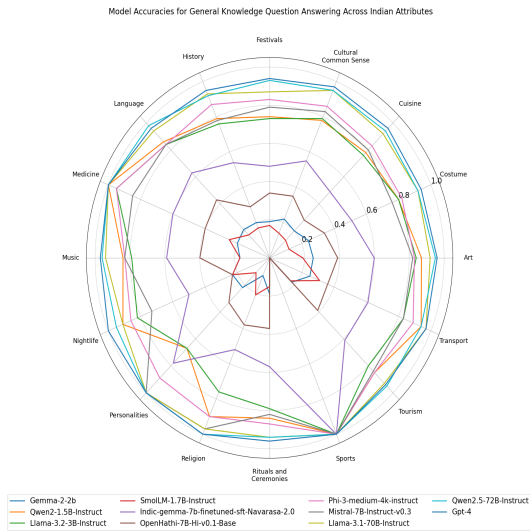


(a) Performance of Language models on Country_Prediction kind of Questions classified on the basis of attributes

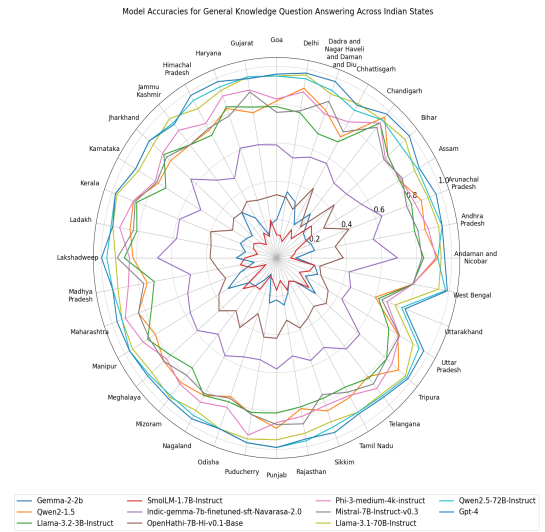


(b) Performance of Language models on Country_Prediction kind of Questions classified on the basis of states

Figure 13: Performance of language models on the data points that falls under country prediction kind of questions in *SANSKRITI* that are further classified on the basis of attributes and states.

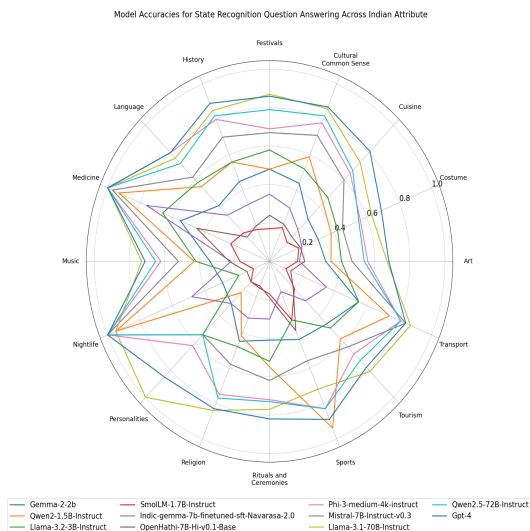


(a) Performance of Language models on General knowledge kind of Questions classified on the basis of attributes.

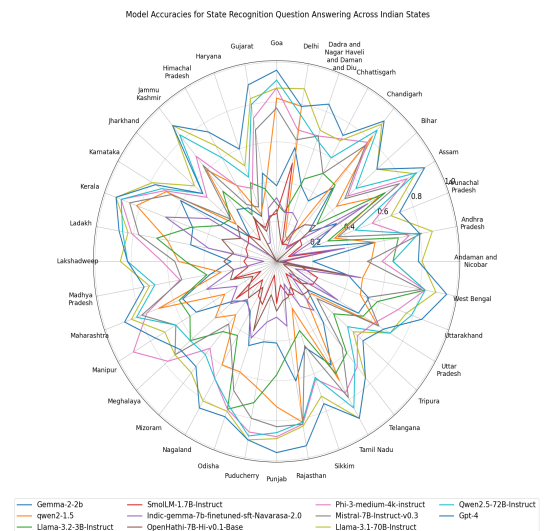


(b) Performance of Language models on General Knowledge kind of Questions classified on the basis of states.

Figure 14: Performance of language models on the data points that falls under General Knowledge kind of questions in *SANSKRITI* that are further classified on the basis of attributes and states.



(a) Performance of Language models on State_Prediction kind of Questions classified on the basis of attributes.



(b) Performance of Language models on State_Prediction kind of Questions classified on the basis of states.

Figure 15: Performance of language models on the data points that falls under State_Prediction kind of questions in *SANSKRITI* that are further classified on the basis of attributes and states.