

The Riddle of Reflection: Evaluating Reasoning and Self-Awareness in Multilingual LLMs Using Indian Riddles

Abhinav P M¹, Ojasva Saxena², Oswald C³, Parameswari Krishnamurthy¹

¹International Institute of Information Technology, Hyderabad

²Independent Researcher

³National Institute of Technology, Tiruchirappalli

abhinav.pm@research.iit.ac.in, ojassaxena03@gmail.com,

oswald.c@nitt.edu, param.krishna@iit.ac.in

Abstract

The extent to which large language models (LLMs) can perform culturally grounded reasoning across non-English languages remains underexplored. This paper examines the reasoning and self-assessment abilities of LLMs across seven major Indian languages- Bengali, Gujarati, Hindi, Kannada, Malayalam, Tamil, and Telugu. We introduce a multilingual riddle dataset combining traditional riddles with context-reconstructed variants and evaluate five LLMs- Gemini 2.5 Pro, Gemini 2.5 Flash, Mistral-Saba, LLaMA-4-Scout, and LLaMA-4-Maverick under seven prompting strategies. In the first stage, we assess riddle-solving performance and find that while Gemini 2.5 Pro performs best overall, few-shot methods yield only marginal gains, and accuracy varies notably across languages. In the second stage, we conduct a self-evaluation experiment to measure reasoning consistency. The results reveal a key finding: a model's initial accuracy is inversely correlated with its ability to identify its own mistakes. Top-performing models such as Gemini 2.5 Pro are overconfident (4.34% True Negative Rate), whereas lower-performing models like LLaMA-4-Scout are substantially more self-aware (42.09% True Negative Rate). These results point to clear gaps in multilingual reasoning and highlight the need for models that not only reason effectively but also recognize their own limitations.

Keywords: cross-lingual reasoning, riddle solving, prompting, Indian languages, self-evaluation

1. Introduction

Large language models (LLMs) such as GPT-4 have demonstrated remarkable capabilities in question-answering and reasoning on a variety of English benchmarks (Achiam et al., 2023). However, this progress has been largely confined to English-centric evaluations, leaving a significant gap in our understanding of their performance in multilingual and cross-cultural contexts (Ghosh et al., 2025), and the same models often fail when asked to apply knowledge in cross-cultural contexts (Liu et al., 2023). Riddles, which often find their origins in traditional folk puzzles, are linguistically rich and context-dependent, making them an ideal stress test for deeper understanding. Solving riddles typically requires multi-step commonsense reasoning, metaphorical or analogical interpretation, and even counterfactual thinking (Lin et al., 2021). As highlighted by the recent survey on puzzle reasoning with LLMs (Giadikiaroglou et al., 2024), models frequently exhibit surface-level pattern matching rather than genuine logical reasoning, particularly when tasks involve creativity, implicit cues, or cultural context. Evaluating large language models (LLMs) on such tasks offers valuable insight into their ability to combine linguistic understanding, creative reasoning, and contextual adaptability.

While existing benchmarks such as RiddleSense

(Lin et al., 2021), BrainTeaser (Jiang et al., 2023), and BiRdQA (Zhang and Wan, 2022) explore riddle solving as a test of commonsense reasoning, they remain limited in their cultural and linguistic scope. This gap highlights the need for methods that can better guide LLMs in culturally-grounded contexts. Prompt engineering offers a powerful approach to this challenge, as studies consistently show that techniques like chain-of-thought and few-shot prompting significantly improve multi-step reasoning (Qiao et al., 2022). The fact that even minimal in-context examples can steer complex behavior (Viswanathan et al., 2024) underscores the value of exemplar-based prompting for equipping LLMs to solve culturally specific riddles. More recently, the RISCORE (Panagiotopoulos et al., 2024) framework introduced the idea of using context-reconstructed exemplars—where new puzzles are automatically generated by altering context but retaining logical structure—to further boost riddle-solving performance in language models.

In this work, we systematically evaluate the riddle-solving and self-assessment abilities of modern large language models across seven major Indian languages—Bengali, Gujarati, Hindi, Kannada, Malayalam, Tamil, and Telugu. While recent work such as (Panagiotopoulos et al., 2024) has examined English riddles, little is known about how models perform when riddles are drawn from multi-

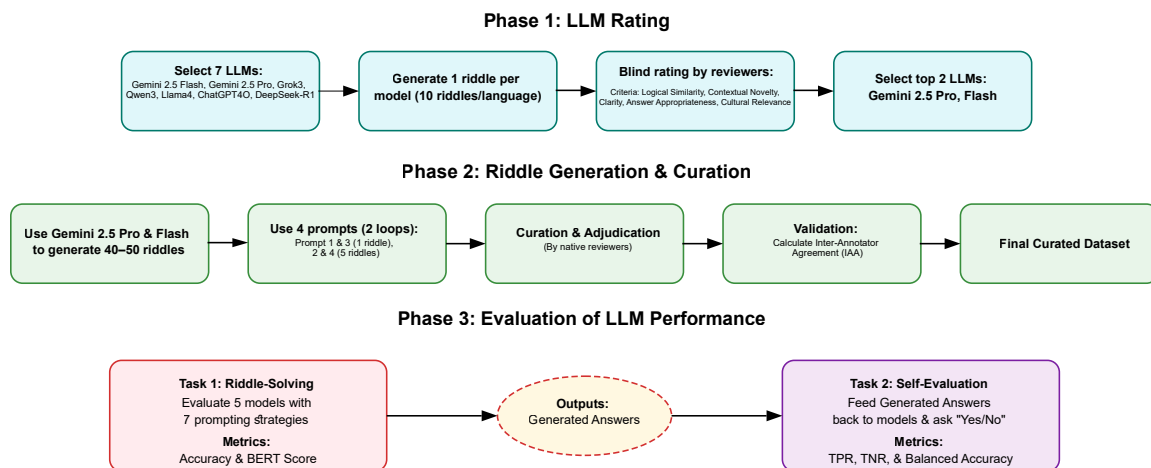


Figure 1: Overview of the comprehensive three-phase methodology, including LLM rating, riddle generation with validation, and a two-stage evaluation of riddle-solving and model self-awareness.

lingual and culturally grounded sources.

Our study seeks to answer the following research questions:

- How effectively can current LLMs solve riddles presented in diverse Indian languages?
- How do different prompting strategies such as zero-shot, few-shot with random, semantically similar, and contextually reconstructed exemplars affect riddle-solving performance across these languages?
- Are LLMs capable of evaluating their own outputs, and what does this reveal about their self-awareness and reasoning consistency?

To answer these questions, we construct a new multilingual riddle dataset¹ that combines traditional riddles with context-reconstructed variants inspired from (Panagiotopoulos et al., 2024). Through extensive experiments with multiple prompting strategies, we analyze how prompting style and language influence model performance in riddle solving. Beyond generation, we further introduce a self-evaluation framework that probes each model’s ability to judge the correctness of its own responses. Together, these investigations provide a comprehensive understanding of how current LLMs perform, generalize from examples, and evaluate their own outputs across diverse Indian languages. An overview of our comprehensive three-phase methodology is presented in Figure 1.

¹Data and code can be found here: https://github.com/ltrc/Riddle_of_reflection_paper

2. Related Work

Multilingual Reasoning in Large Language Models. Recent work has highlighted significant challenges in multilingual reasoning for large language models. Recent work demonstrates that cross-lingual reasoning via test-time scaling while scaling inference-time compute improves reasoning performance (Yong et al., 2025), a persistent gap remains between English and non-English languages, with models showing degraded performance in low-resource languages. This finding is further corroborated by (Schut et al., 2025), who investigate whether multilingual LLMs think in English. Their analysis reveals that even when prompted in other languages, many multilingual models perform internal reasoning in English before translating back to the target language, leading to cultural and contextual misalignments. These findings underscore the importance of evaluating models on culturally-grounded tasks in their native languages. Many recent benchmark studies confirm that LLMs still struggle with multilingual and culturally nuanced reasoning. For example, benchmarks like IndicGenBench (Singh et al., 2024) show a notable performance gap between English and a variety of Indian languages, with large drops in low-resource settings. Similarly, the MultiNRC (Fabbri et al., 2025) benchmark reveals that when models are tasked with native, culturally grounded reasoning problems (including word-play, riddles, and tradition-based puzzles) in non-English languages, none of the state-of-the-art models exceeded 50% accuracy.

Prompting Strategies and Reasoning in LLMs. Prompting has emerged as one of the most effective approaches for enhancing reasoning at in-

ference time. (Qiao et al., 2022) provide a comprehensive survey on reasoning with language model prompting, identifying techniques including few-shot, chain-of-thought, and self-consistency prompting, showing that providing exemplars can significantly improve multi-step reasoning. Quest-Bench (Li et al., 2025) assessed whether LLMs can formulate clarifying questions to improve their reasoning-demonstrating that effective reasoning often involves identifying missing contextual information. The challenges of puzzle-solving for LLMs have been explored through various task-specific studies. (Tyagi et al., 2024) investigate grid puzzle solving and identify specific failure modes where LLMs struggle with spatial reasoning and multi-step constraint satisfaction. Their analysis shows that even step-by-step reasoning can falter when tasks require maintaining complex state representations. The RISCORE framework by (Panagiotopoulos et al., 2024) introduce context-reconstructed exemplars (i.e., same logical pattern, altered context) as an effective way to stimulate reasoning in riddles. Their approach demonstrates improvements in English riddle-solving, but does not extend to multilingual or culturally diverse settings.

Self-Evaluation and Model Self-Awareness. The ability of LLMs to evaluate their own outputs has emerged as a critical dimension of model reliability. (Huang et al., 2024) introduce a framework for self-evaluation based on glass-box features, analyzing internal model states to predict output quality. (Ren et al., 2023) show that self-evaluation can improve selective generation, allowing models to abstain from answering when uncertain. However, (Liu et al., 2024) reveal significant limitations through their Think-Solve-Verify framework, finding that while models can sometimes identify errors in their reasoning chains, this self-awareness is inconsistent and often fails in complex reasoning tasks. Most recently, (Nguyen et al., 2025) demonstrate that LLMs exhibit evaluation awareness-the ability to recognize when they are being evaluated-and that this awareness can be both probed and steered.

3. Methodology and Experiment

The methodology of this study is structured into three main phases: (1) LLM selection for riddle generation through a comprehensive rating process, (2) generation and curation of a novel dataset of context-reconstructed riddles using the selected LLMs, and (3) a two-stage evaluation of LLM performance, assessing both their riddle-solving capabilities and their capacity for self-evaluation. A contextually reconstructed riddle, inspired by the RISCORE framework (Panagiotopoulos et al.,

2024), is defined as a riddle that maintains the logical structure of an original riddle while completely altering its theme or context. An example illustrating this transformation is shown in Figure 2. An outline of the overall workflow is illustrated in Figure 1.

Original Riddle

हिंदी: ऐसी कौन - सी जगह है, जहाँ पर सड़क है पर गाड़ी नहीं, जंगल है पर पेड़ नहीं और शहर है पर घर नहीं ?
उत्तर: नक्शा।

English Translation: What is a place that has roads but no cars, forests but no trees, and cities but no houses? Answer: Map.

Context-Reconstructed Riddle 1

हिंदी: ऐसी कौन सी चीज़ है, जिसमें लोग हैं पर जान नहीं, जगह है पर माप नहीं, और समय है पर गति नहीं?
उत्तर: तस्वीर।

English Translation: What is a thing that has people but no life, space but no measurement, and time but no movement? Answer: Photograph.

Context-Reconstructed Riddle 2

हिंदी: ऐसी कौन है, इसमें दुनिया बसी है पर कोई रहता नहीं, बातें होती हैं पर जुबान नहीं, चीज़ें दिखती हैं पर पकड़ नहीं सकते।
उत्तर: मोबाइल स्क्रीन।

English Translation: What has the world inside it, but no one lives there; conversations happen, but there's no tongue? Answer: Mobile Screen

Figure 2: An original Hindi riddle (top) and two of its context-reconstructed variants. Note how the core reasoning pattern is maintained while the theme and answer are altered.

3.1. Phase 1: LLM Selection

To identify suitable LLMs for generating context-reconstructed riddles, seven trending models are evaluated: Gemini 2.5 Pro, Gemini 2.5 Flash, Grok 3, Qwen3-235B-A22B (Yang et al., 2025), LLaMA 4 Maverick 17B Instruct (128E), ChatGPT-4o, and DeepSeek-R1 (Guo et al., 2025). These models are selected based on their reported performance on various natural language tasks and their potential to handle multilingual contexts. For each model, 10 original riddles in each of the seven Indian languages (Bengali, Gujarati, Hindi, Kannada, Malayalam, Tamil, and Telugu) are provided, along with their answers. The models are prompted to generate one context-reconstructed riddle per original riddle, maintaining the same logical pattern while altering the context.

The generated riddles are then assessed by native-speaking expert reviewers in a blind rating process. Each reconstruction is scored against five key criteria: (1) *Logical Similarity*, the preservation of the original reasoning pattern; (2) *Contextual Novelty*, the creativity and distinctiveness of the new context; (3) *Clarity and Coherence*, the linguistic clarity and structural coherence of the riddle; (4) *Answer Appropriateness*, the suitability of the new answer; and (5) *Cultural Relevance*, the alignment with cultural norms and idioms of the target language. The overall LLM ratings across the seven Indian languages are presented in Figure 3. The evaluation shows that Gemini 2.5 Pro and Gemini 2.5 Flash perform best among the tested models, demonstrating a stronger ability to capture linguistics-

tic nuances and cultural aspects specific to Indian languages.




Figure 3: LLM Performance in Reconstructed Contextual Riddle Generation Across Indian Languages.

3.2. Phase 2: Generation and Curation of Reconstructed-Context Riddles

Using the selected Gemini 2.5 Pro and Gemini 2.5 Flash models, we generated a large collection of context-reconstructed riddles. The source material consists of 96 traditional, publicly sourced riddles for each of the seven Indian languages. To ensure a diverse pool of candidate context reconstructions, we use four distinct prompt variations, ranging from simple requests to more detailed, analytical instructions. Each prompt is executed twice, producing a pool of 40–50 candidate reconstructions for each of the 96 original riddles per language. From this pool, annotators manually review the candidates and select the top four reconstructed riddles per original riddle.

Curation and Inter-Annotator Agreement: To ensure the validity and reliability of the final dataset, all generated candidates undergo a careful curation and validation process. For each language, a team of native-speaking annotators (three for five languages and two for Malayalam and Bengali)² evaluates the reconstructions to identify those of the highest quality. Each annotator marks “Y” for acceptable candidates and assigns a quality score from 0 to 10, reflecting an overall assessment based on the same five criteria used during the initial LLM rating phase. Since only top-quality riddles are considered, most scores fall within the 7–10 range.

To measure the consistency of these subjective judgments, we conduct an inter-annotator agreement (IAA) study by calculating Cohen’s Kappa for

²The annotation team for Malayalam and Bengali consists of two members due to the limited availability of native expert reviewers.

each annotator pair. The average IAA across all languages ranges from 55% to 61%. This moderate level of agreement is expected, as evaluating criteria like *Contextual Novelty* and *Cultural Relevance* involves significant subjective interpretation. Annotators often differed in their opinions on how creative a reconstruction was or how natural its phrasing sounded, which led to variations in the scores. This finding confirms that selecting the “best” creative contextual reconstructions is an inherently subjective task. In cases of disagreement, annotators discuss collaboratively to finalize the top four riddles for inclusion in the curated dataset.

3.3. Phase 3: Evaluation of LLM Performance

LLM performance is evaluated in two distinct stages: first, by assessing their direct riddle-solving capabilities, and second, by probing their ability to perform self-evaluation.

3.3.1. Riddle-Solving Task

The riddle-solving capabilities of LLMs are assessed using five models selected for their strong performance and general availability within the research community: Gemini 2.5 Pro, Gemini 2.5 Flash, LLaMA 4 Maverick 17B Instruct (128E), LLaMA 4 Scout 17B Instruct (16E), and Mistral-Saba. The 96 original riddles in each language serve as the test set. Seven distinct prompting strategies are used to comprehensively evaluate the impact of exemplar-based reasoning. The exact prompt templates used for these experiments are provided in Appendix A.2.

- **Zero-Shot:** The model receives only the test riddle.
- **2-Shot and 4-Shot Random:** The prompt includes two or four randomly selected riddles of the same language from the curated dataset as exemplars.
- **2-Shot and 4-Shot Semantic Similarity:** Exemplars are the two or four riddles of the same language from the dataset most semantically similar to the test riddle, as determined using LaBSE (Feng et al., 2020) embeddings and cosine similarity.
- **2-Shot and 4-Shot Contextual Reconstructed:** Two or four context-reconstructed riddles derived from the original riddle of the same language are used as exemplars, following the (Panagiotopoulos et al., 2024) inspired approach.

The models’ outputs are evaluated against the ground truth using two primary metrics:

- **Accuracy:** The percentage of generated answers matching the ground truth.
- **BERTScore:** The F1 score of semantic similarity, computed using the 'bert-base-multilingual-cased' (Devlin et al., 2019) model to handle partially correct or varied answers.

3.3.2. Self-Evaluation Task

To examine the models' self-awareness and reasoning consistency, we design a second experiment. In this task, the answers previously generated by each model under the seven prompting strategies for each language are used. The same model is then presented with the original riddle and its own candidate answer and asked to perform a binary classification to determine whether its generated answer correctly solves the riddle. The model is instructed to respond with a single word, "Yes" or "No," indicating whether the answer is correct.

This experiment measures two key aspects of self-awareness, defined using a standard confusion matrix:

- **True Positive Rate (TPR):** The model's ability to correctly identify its own correct answers (i.e., saying "Yes" when its initial answer is right). This measures confidence and consistency.
- **True Negative Rate (TNR):** The model's ability to correctly identify its own incorrect answers (i.e., saying "No" when its initial answer was wrong). This measures the capacity for self-correction and awareness of failure.

These two metrics, along with the derived Balanced Accuracy (the average of TPR and TNR), provide a deeper understanding of model reliability beyond simple generation accuracy, revealing critical insights into overconfidence and reasoning stability.

4. Result and Discussion

This section presents the results from both experiments of the evaluation: (1) riddle-solving performance across five large language models (LLMs) and seven prompting strategies, and (2) an analysis of model self-awareness based on their ability to evaluate their own answers. Together, they highlight the models' reasoning performance, ability to generalize from examples, and consistency in self-assessment across languages.

4.1. Riddle-Solving Performance

The initial evaluation focused on the models' ability to generate correct answers for 96 original riddles

per language under seven different prompting conditions. Five LLMs: Gemini 2.5 Pro (G. Pro), Gemini 2.5 Flash (G. Flash), Mistral-Saba (M. Saba), LLaMA-4-Scout (L. Scout), and LLaMA-4-Maverick (L. Mav) are evaluated on riddles across seven Indian languages: Bengali, Gujarati, Hindi, Kannada, Malayalam, Tamil, and Telugu. Model performance is measured using two metrics: *Accuracy* and *BERTScore F1*. Detailed numerical results for all combinations are presented in Table 1.

4.1.1. Overall Model Performance

Gemini 2.5 Pro consistently outperformed other models across most languages and prompting settings. As shown in Figure 4 and Figure 5, it achieves the highest accuracy in riddle-solving, making it the best-performing model in our evaluation. Its highest accuracies are observed in Gujarati (63.54% with 2-shot random) and Tamil (56.25% in zero-shot), while its BERTScore F1 peaks at 93.04 in Gujarati (4-shot random). Gemini 2.5 Flash records the second-highest performance, while Mistral-Saba, LLaMA 4 Maverick, and LLaMA 4 Scout form a lower-performing group.

Another observation from both plots is the significant and persistent gap between Accuracy (solid lines) and BERTScore F1 (dashed lines). For all models, BERTScore F1 values are consistently high, typically ranging from 75% to 93%. In contrast, accuracy is substantially lower, with Gemini 2.5 Pro averaging around 40% and other models performing below 30%. This wide gap suggests that while the models are proficient at generating semantically relevant or contextually appropriate answers, they frequently fail to identify the precise, creative solution required by riddles. Even with this wide gap, the BERTScore F1 curves largely follow the same trend as the accuracy curves. This indicates that the factors affecting a model's ability to produce the exact answer also influence its ability to produce a semantically close one.

4.1.2. Impact of Prompting Strategies

Model performance across the seven prompting strategies, averaged over all languages, is shown in Figure 4. Overall, the results show that the choice of prompting strategy has only a modest impact on performance. Few-shot approaches such as 4-shot Random (average accuracy of 23.24%) slightly outperform the zero-shot (21.47%), but none of the prompting strategies provides a clear or consistent advantage. Gemini 2.5 Pro, for example, achieves its peak average accuracy with 4-shot Random (42.11%), but its performance is nearly identical in the zero-shot setting (41.52%). The context-reconstructed examples also do not offer a consistent advantage, suggesting that simply in-

Language	P. Strategy	G. Pro		G. Flash		M. Saba		L. Scout		L. Mav	
		Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
Bengali	0-shot Std.	41.67	88.21	23.96	84.72	13.02	80.97	10.42	80.60	11.46	81.26
	2-shot Rand.	39.58	87.27	22.92	84.74	17.71	83.08	6.25	79.27	11.46	81.03
	4-shot Rand.	41.67	88.50	28.12	84.62	17.71	82.90	9.38	80.42	13.54	80.76
	2-shot Sem.	42.71	87.35	28.12	84.06	15.62	82.67	8.33	80.96	13.54	80.89
	4-shot Sem.	37.50	86.19	25.00	82.91	16.67	82.41	11.46	81.64	14.58	80.34
	2-shot Ctx.	42.71	87.19	29.17	84.35	19.79	83.53	8.33	80.38	11.46	80.90
	4-shot Ctx.	43.75	86.26	26.04	84.31	19.79	83.58	7.29	80.48	15.62	82.13
Gujarati	0-shot Std.	57.29	92.52	42.71	88.54	28.12	84.33	21.88	82.51	23.96	82.05
	2-shot Rand.	63.54	92.68	41.67	88.40	32.29	85.36	21.88	83.08	30.21	84.62
	4-shot Rand.	62.50	93.04	42.71	88.81	33.33	86.73	26.04	83.94	30.21	83.76
	2-shot Sem.	61.46	91.74	43.75	88.81	28.12	84.06	20.83	83.69	27.08	83.52
	4-shot Sem.	50.00	89.24	44.79	89.28	28.12	84.40	18.75	81.97	20.83	81.68
	2-shot Ctx.	56.25	90.41	43.75	89.37	26.04	84.50	27.08	83.49	28.12	82.84
	4-shot Ctx.	48.96	88.11	43.75	88.36	31.25	86.05	23.96	83.21	29.17	84.09
Hindi	0-shot Std.	45.83	90.12	36.46	87.72	37.50	87.50	11.46	81.64	27.08	84.37
	2-shot Rand.	43.75	89.45	37.50	87.71	45.83	89.54	12.50	81.78	20.83	82.46
	4-shot Rand.	46.88	90.45	36.46	87.88	38.54	88.63	17.71	83.04	22.92	82.90
	2-shot Sem.	48.96	89.11	39.58	88.33	37.50	87.90	20.83	83.63	22.92	83.17
	4-shot Sem.	41.67	86.95	35.42	87.26	34.38	88.29	21.88	83.74	26.04	84.60
	2-shot Ctx.	45.83	88.35	38.54	88.34	39.58	88.38	16.67	81.69	26.04	84.14
	4-shot Ctx.	41.67	86.63	38.54	87.81	40.62	89.37	16.67	83.24	26.04	83.88
Kannada	0-shot Std.	20.83	83.80	19.79	82.49	7.29	78.45	5.21	77.77	7.29	77.76
	2-shot Rand.	23.96	83.58	18.75	82.07	12.50	79.76	3.12	76.47	11.46	78.49
	4-shot Rand.	29.17	84.29	20.83	83.15	10.42	79.49	3.12	77.18	10.42	77.42
	2-shot Sem.	23.96	80.98	17.71	81.80	10.42	79.40	7.29	78.80	8.33	78.49
	4-shot Sem.	21.88	79.38	15.62	81.52	12.50	79.34	7.29	78.39	8.33	77.18
	2-shot Ctx.	21.88	81.08	12.50	80.07	10.42	78.48	5.21	77.48	7.29	77.28
	4-shot Ctx.	22.92	80.77	14.58	80.91	8.33	78.56	4.17	76.84	9.38	78.22
Malayalam	0-shot Std.	42.71	88.47	23.96	82.85	2.08	75.30	4.17	75.83	6.25	77.30
	2-shot Rand.	36.46	87.04	22.92	82.40	1.04	76.38	3.12	75.02	10.42	78.55
	4-shot Rand.	41.67	87.59	20.83	82.27	6.25	76.88	2.08	75.06	8.33	78.08
	2-shot Sem.	45.83	87.19	23.96	82.87	3.12	76.50	6.25	76.98	8.33	77.23
	4-shot Sem.	41.67	86.48	20.83	82.30	6.25	77.36	5.21	76.57	10.42	78.31
	2-shot Ctx.	32.29	83.31	19.79	81.72	3.12	76.24	6.25	76.72	9.38	77.42
	4-shot Ctx.	43.75	86.72	15.62	80.50	2.08	76.20	7.29	79.25	10.42	77.98
Tamil	0-shot Std.	56.25	91.58	31.25	85.63	15.62	81.40	9.38	79.57	16.67	81.56
	2-shot Rand.	52.08	90.60	35.42	86.26	28.12	84.76	9.38	79.24	19.79	82.75
	4-shot Rand.	52.08	90.72	36.46	86.79	25.00	83.40	7.29	77.94	17.71	82.09
	2-shot Sem.	50.00	89.34	32.29	85.96	21.88	83.32	9.38	79.45	19.79	82.28
	4-shot Sem.	41.67	87.12	33.33	86.36	22.92	83.65	12.50	80.26	18.75	81.68
	2-shot Ctx.	52.08	89.28	32.29	85.05	19.79	82.38	4.17	78.23	18.75	82.74
	4-shot Ctx.	45.83	87.73	35.42	86.73	21.88	83.07	11.46	79.67	21.88	82.31
Telugu	0-shot Std.	26.04	86.19	12.50	81.92	3.12	77.78	4.17	77.08	4.17	77.37
	2-shot Rand.	23.96	85.32	15.62	81.98	10.42	79.63	3.12	75.53	6.25	77.40
	4-shot Rand.	20.83	84.51	15.62	82.00	7.29	79.18	4.17	76.79	6.25	77.73
	2-shot Sem.	17.71	81.14	12.50	81.28	8.33	78.85	4.17	77.38	7.29	78.56
	4-shot Sem.	21.88	82.02	12.50	81.90	8.33	80.35	2.08	77.19	6.25	78.27
	2-shot Ctx.	25.00	83.37	9.38	80.71	5.21	79.27	5.21	78.22	6.25	77.72
	4-shot Ctx.	21.88	82.02	11.46	81.85	13.54	80.39	3.12	77.32	6.25	78.90

Table 1: Detailed Model Performance Across Languages and Prompt Strategies. Acc denotes Accuracy (%), F1 denotes BERTScore F1. Highest scores per model within each language are bolded.

cluding examples with similar reasoning patterns is not enough to address the core challenges of this task across different languages. The overall flatness of the accuracy lines suggests that prompting style has limited influence compared to the model's inherent strength in this type of task.

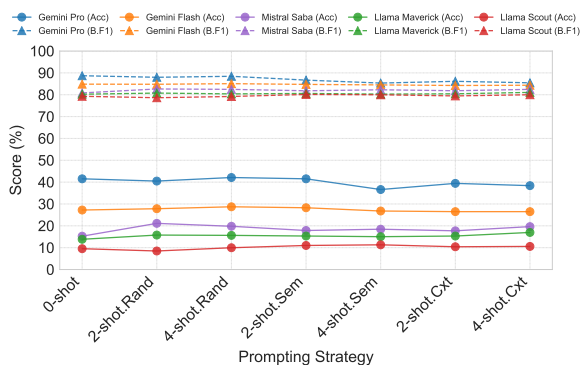


Figure 4: Average riddle-solving accuracy (%) and BERTScore F1 across seven prompting strategies for five LLMs. The plot illustrates the minor variation in model performance based on the prompting method.

4.1.3. Performance Disparities Across Indian Languages

Performance varied significantly across languages as shown in Figure 5. The plot shows a clear trend where all models perform significantly better on riddles in Gujarati, Hindi, and Tamil. Gemini 2.5 Pro, for instance, achieves its peak average accuracy in Gujarati (57.14%) and Tamil (50.00%). Conversely, performance drops dramatically for all

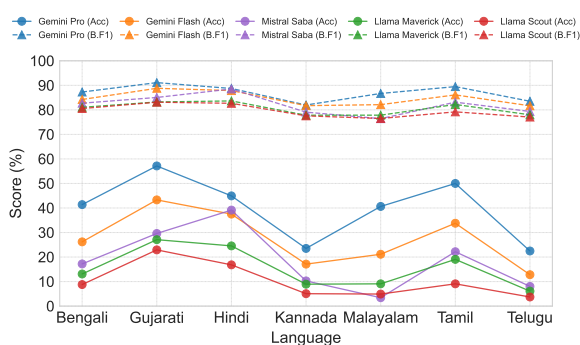


Figure 5: Average riddle-solving accuracy (%) and BERTScore F1 across seven Indian languages for five LLMs. The plot illustrates the variation in model performance based on the target language.

models in Kannada, Malayalam, and particularly Telugu. In Malayalam, for example, Mistral-Saba's average accuracy falls to just 3.42%, and even

the top-performing Gemini 2.5 Pro only reaches 40.63%.

This suggests that model proficiency is not uniform across all Indian languages, likely reflecting disparities in training data and cultural context representation. Despite the variation in accuracy, BERTScore F1 values remain consistently high across all settings, indicating that models often produce semantically relevant responses even when the exact answer is incorrect.

4.2. Performance on the Self-Evaluation Task

While initial accuracy measures a model's ability to generate a correct answer, it does not reveal the stability or reliability of its reasoning. To probe this deeper layer of cognition, we conducted a self-evaluation task in which each model assesses its own previously generated answers using a simple "Yes" or "No" response. Table 2 reports the *True Positive Rate* (TPR), *True Negative Rate* (TNR), and the derived *Balanced Self-Evaluation Accuracy*.

4.2.1. The Overconfidence of High Performers

The most significant finding from this task is the strong overconfidence shown by the top-performing models. Gemini 2.5 Pro, despite having the highest initial generation accuracy (39.09%), shows an almost complete inability to recognize its own errors. Its True Positive Rate (TPR) was a perfect 100%, meaning it correctly confirmed every single one of its right answers. However, its True Negative Rate (TNR) is only 4.34%, indicating that when it generates an incorrect answer, it still confidently claims it is correct more than 95% of the time. This pattern of high TPR and very low TNR is also observed in Gemini 2.5 Flash and LLaMA-4-Maverick.

4.2.2. The Self-Awareness of Weaker Performers

In contrast, models with lower initial generation accuracy show significantly better self-awareness. LLaMA-4-Scout, which has the lowest initial accuracy (9.95%), achieves the highest TNR (42.09%). This means it can correctly identify and reject its own incorrect answers far more effectively than any other model. Mistral-Saba shows a similar trend, combining modest initial accuracy with a strong TNR (43.57%). This suggests that these models, while less capable at generation, possess a more reliable internal mechanism for evaluating the validity of an answer.

Model	Avg. Initial Acc. (%)	Self-Eval. (TPR)(%)	Self-Eval. (TNR)(%)	Balanced Self-Eval. Acc. (%)
Gemini-2.5-Pro	39.09	100.00	4.34	52.17
Gemini-2.5-Flash	26.68	99.58	7.66	53.62
LLaMA-4-Maverick	15.09	96.60	10.40	53.50
LLaMA-4-Scout	9.95	95.73	42.09	68.91
Mistral-Saba	18.21	86.28	43.57	64.93

Table 2: Overall Model Performance on Riddle Generation and Self-Evaluation Accuracy across seven Indian languages and prompting strategies. Best values per column are bolded.

4.2.3. The Performance vs. Self-Awareness Trade-off

The quadrant visualization in Figure 6 maps self-awareness against initial generation performance trade-off. The x-axis represents generation performance, while the y-axis represents self-awareness as measured by Balanced Self-Evaluation Accuracy (the average of TPR and TNR). The models fall into distinct categories:

- **Overconfident Performers (Bottom-Right):** Gemini 2.5 Pro and Flash are high-performing but lack self-awareness.
- **Cautious and Aware (Top-Left):** LLaMA-4-Scout and Mistral-Saba are weaker performers but are more capable of identifying their own mistakes.
- **Unreliable (Bottom-Left):** LLaMA-4-Maverick struggled in both generation and self-evaluation.

Crucially, no model falls into the ideal top-right quadrant of high performance and high self-awareness. This inverse relationship between initial accuracy and the ability to recognize errors is a key finding, indicating that reasoning confidence and self-awareness may not yet co-develop in current LLMs.

4.3. Discussion

The results of our study reveal three key insights: (1) **Culturally grounded reasoning remains challenging:** Even the top-performing models struggle to consistently produce correct answers, indicating that while LLMs can generate semantically relevant responses, they often miss the precise, creative insights required for riddle-solving. (2) **Limited Gains from Few-Shot Prompting:** Adding random, contextual, or semantically similar exemplars provides only modest and inconsistent improvements, suggesting that for tasks requiring a blend of creative reasoning and deep cultural knowledge. (3) **Performance and self-awareness**

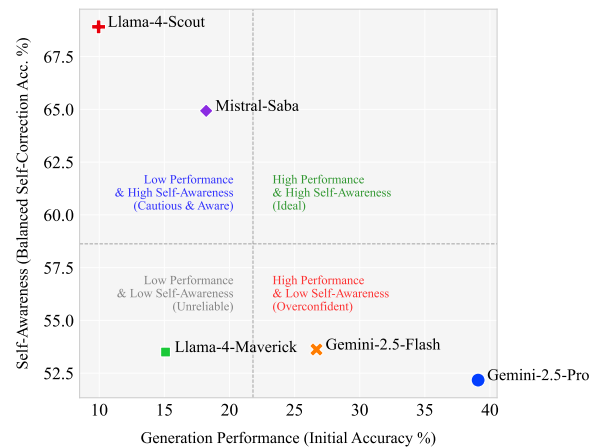


Figure 6: LLM performance vs. self-awareness in riddle solving. The x-axis represents initial riddle-solving accuracy, and the y-axis shows balanced self-evaluation accuracy.

are inversely related: High-performing models, such as Gemini 2.5 Pro, show near-perfect confidence but often fail to recognize their own mistakes, while lower-performing models like LLaMA 4 Scout are more cautious and better at evaluating their answers.

Overall, these results highlight persistent gaps in multilingual reasoning and emphasize the need for culturally grounded, reflective evaluation frameworks across languages. Our objective is to better understand the limitations of multilingual reasoning and model self-evaluation, rather than to propose mitigation strategies or architectural solutions. Through a controlled evaluation framework and empirical observations across languages and prompting strategies, we aim to provide insights that can inform future work on more culturally grounded evaluation of large language models.

5. Conclusion

This paper presented a systematic evaluation of large language models' reasoning and self-

assessment capabilities across seven major Indian languages using riddle-solving as a culturally grounded task. Our three-phase study-covering model selection, dataset curation, and dual-stage evaluation-showed that while models like Gemini 2.5 Pro achieved the highest accuracy, few-shot prompting offered only limited gains, and no single strategy proved consistently effective. Performance varied widely across languages, underscoring the ongoing challenges of multilingual reasoning. The self-evaluation results further reveal that stronger models tend to be overconfident, whereas weaker ones show greater awareness of their own mistakes. This inverse relationship suggests that current LLMs can generate likely answers but still struggle to check whether those answers are actually correct. Future work will expand this analysis to additional languages and reasoning types while exploring methods to enhance model self-awareness through reflective training, and will also investigate the development of a semi-automated pipeline for dataset generation and curation to improve scalability while maintaining linguistic and cultural validity.

6. Ethics Statement

The development of foundational LLMs has been overwhelmingly English-centric, a trend that risks creating a significant digital divide and marginalizing non-English native speakers. In a nation like India, where estimates suggest only 10-15% of the population understands English, an over-reliance on Anglocentric models is an ethical concern.

For LLMs to be inclusive and beneficial tools, they must be able to grasp the cultural context, nuance, and implicit cues of diverse linguistic communities. This capability is essential for linguistic minorities and marginalized communities to feel represented and to interact safely with AI systems, rather than encountering "cultural and contextual misalignments".

This work aims to foster the development of foundational models that are not just multilingual, but truly multicultural in their reasoning capabilities, ensuring that the benefits of AI are accessible and relevant to all.

7. Acknowledgements

We sincerely thank Professor Dipti Misra Sharma for her insightful feedback, particularly the suggestion to conduct an Inter-Annotator Agreement study. We also extend our gratitude to R S Mughil Srinivasan and Kesavan T for their contributions to the Tamil riddle collection. Our thanks also go to the annotation teams across the seven Indian

languages for their effort and dedication in curating and validating the multilingual riddle dataset.

8. Bibliographical References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Alexander R Fabbri, Diego Mares, Jorge Flores, Meher Mankikar, Ernesto Hernandez, Dean Lee, Bing Liu, and Chen Xing. 2025. Multinrc: A challenging and native multilingual reasoning evaluation benchmark for llms. *arXiv preprint arXiv:2507.17476*.

Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. Language-agnostic bert sentence embedding. *arXiv preprint arXiv:2007.01852*.

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*.

Panagiotis Giadikiaroglou, Maria Lymperaiou, Giorgos Filandrianos, and Giorgos Stamou. 2024. Puzzle solving using reasoning of large language models: A survey. *arXiv preprint arXiv:2402.11291*.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

Hui Huang, Yingqi Qu, Jing Liu, Muyun Yang, Bing Xu, Tiejun Zhao, and Wenpeng Lu. 2024. Self-evaluation of large language model based on glass-box features. *arXiv preprint arXiv:2403.04222*.

- Yifan Jiang, Filip Ilievski, Kaixin Ma, and Zhivar Sourati. 2023. [BRAINTEASER: Lateral thinking puzzles for large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14317–14332, Singapore. Association for Computational Linguistics.
- Belinda Z Li, Been Kim, and Zi Wang. 2025. Quest-bench: Can llms ask the right question to acquire information in reasoning tasks? *arXiv preprint arXiv:2503.22674*.
- Bill Yuchen Lin, Ziyi Wu, Yichi Yang, Dong-Ho Lee, and Xiang Ren. 2021. Riddlesense: Reasoning about riddle questions featuring linguistic creativity and commonsense knowledge. *arXiv preprint arXiv:2101.00376*.
- Chen Cecilia Liu, Fajri Koto, Timothy Baldwin, and Iryna Gurevych. 2023. Are multilingual llms culturally-diverse reasoners? an investigation into multicultural proverbs and sayings. *arXiv preprint arXiv:2309.08591*.
- Zhendong Liu, Changhong Xia, Wei He, and Chongjun Wang. 2024. Trustworthiness and self-awareness in large language models: An exploration through the think-solve-verify framework. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 16855–16866.
- Jord Nguyen, Khiem Hoang, Carlo Leonardo Atubato, and Felix Hofstätter. 2025. Probing evaluation awareness of language models. *arXiv preprint arXiv:2507.01786*.
- Ioannis Panagiotopoulos, Giorgos Filandrianos, Maria Lymperaiou, and Giorgos Stamou. 2024. Riscore: Enhancing in-context riddle solving in language models through context-reconstructed example augmentation. *arXiv preprint arXiv:2409.16383*.
- Shuofei Qiao, Yixin Ou, Ningyu Zhang, Xiang Chen, Yunzhi Yao, Shumin Deng, Chuanqi Tan, Fei Huang, and Huajun Chen. 2022. Reasoning with language model prompting: A survey. *arXiv preprint arXiv:2212.09597*.
- Jie Ren, Yao Zhao, Tu Vu, Peter J Liu, and Balaji Lakshminarayanan. 2023. Self-evaluation improves selective generation in large language models. In *Proceedings on*, pages 49–64. PMLR.
- Lisa Schut, Yarin Gal, and Sebastian Farquhar. 2025. Do multilingual llms think in english? *arXiv preprint arXiv:2502.15603*.
- 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*.
- Nemika Tyagi, Mihir Parmar, Mohith Kulkarni, Aswin Rrv, Nisarg Patel, Mutsumi Nakamura, Arindam Mitra, and Chitta Baral. 2024. Step-by-step reasoning to solve grid puzzles: Where do llms falter? *arXiv preprint arXiv:2407.14790*.
- Vijay Viswanathan, Kiril Gashteovski, Kiril Gashteovski, Carolin Lawrence, Tongshuang Wu, and Graham Neubig. 2024. Large language models enable few-shot clustering. *Transactions of the Association for Computational Linguistics*, 12:321–333.
- An Yang, Anpeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Zheng-Xin Yong, M Farid Adilazuarda, Jonibek Mansurov, Ruochen Zhang, Niklas Muennighoff, Carsten Eickhoff, Genta Indra Winata, Julia Kreutzer, Stephen H Bach, and Alham Fikri Aji. 2025. Crosslingual reasoning through test-time scaling. *arXiv preprint arXiv:2505.05408*.
- Yunxiang Zhang and Xiaojun Wan. 2022. Birdqa: A bilingual dataset for question answering on tricky riddles. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11748–11756.

A. Appendix

This appendix provides supplementary details, including prompt templates, examples, and detailed results by language and model.

A.1. Dataset and Annotation Notes

Each annotator evaluated roughly 4,200 generated candidate riddles per language. Across the seven languages, this resulted in approximately 29,500 reconstructed riddles being reviewed in total. This manual review helped ensure the quality of the reconstructed riddles and reduce noise in the final dataset.

A.2. Prompt Templates

The same prompt templates are used across all languages for zero-shot and few-shot prompting. Representative templates are shown in Figures 7, 8, 9, 10 and 11.

Prompt : Zero-Shot

You are an expert in solving Tamil riddles.
 Your task is to analyze the following Tamil riddle.
 Provide ONLY the answer. Do not include any explanations, introductions, or any text other than the answer itself.

Tamil Riddle :
 பச்சை பொட்டிக்குள் வெள்ளை முத்துகள்?

Answer :

Figure 7: 0-shot prompt (Tamil)

Prompt : 2-shot Random

You are an expert in solving Hindi riddles.
 Below are some examples of Hindi riddles and their answers.
 Study these examples carefully.
 For the final riddle presented, provide ONLY the answer. Do not include any explanations, introductions, or any text other than the answer itself.

Examples:
 Example 1:
 Riddle: चुबान मुझे मार डालती है, पर भरोसा मुझे जिंदा रखता है। बताओ क्या?
 Answer: राज़

Example 2:
 Riddle: कौन सी चीज़ है जिसकी आवाज़ है पर चुबान नहीं?
 Answer: घंटी

Now, solve this riddle:
 Riddle : वह क्या चीज़ है जिसे आदमी छुपा कर और औरत दिखा कर चलती है ?
 Answer :

Figure 8: 2-shot random prompt (Hindi). Exemplars are randomly drawn from the full Hindi riddle dataset.

- **Zero-shot prompt.** The model is instructed to act as an expert riddle solver for the target language and provide only the final answer. See 7 for an example zero-shot prompt in Tamil.
- **Random few-shot.** Examples are randomly selected from the training pool for the target language and presented as exemplars. See

Prompt : 2-shot Semantically Similar

You are an expert in solving Malayalam riddles.
 Below are some examples of Malayalam riddles and their answers.
 These examples were chosen because they are semantically similar to the kind of riddle you need to solve, even though their specific reasoning patterns might vary.
 Study these examples carefully.
 For the final riddle presented, provide ONLY the answer. Do not include any explanations, introductions, or any text other than the answer itself.

Examples:
 Example 1:
 Riddle: അടിയൻ്റെ മായ്ക്കുന്നവൻ, നടുക്ക് മരം, മീനെ മുർച്ച?
 Answer: പെൻസിൽ

Example 2:
 Riddle: കുരിശിൽ ഒരു തരി പൊന്ന്?
 Answer: മിന്നാമിനുങ്ങ്

Now, solve this riddle:
 Riddle : ചെടി ചെടിയന്മേൽകായ് കയ്യിന്മേൽ ചെടി?
 Answer :

Figure 9: 2-shot semantically similar prompt (Malayalam). Exemplars are based on semantic similarity to the target riddle, drawn from the full Malayalam riddle dataset.

Figure 8 for an example 2-shot random prompt in Hindi.

- **Semantically similar few-shot.** Exemplars are selected based on surface semantic similarity to the target riddle. See Figure 9 for a 2-shot semantically similar prompt in Malayalam.
- **Context reconstructed few-shot.** Exemplars were manually selected from the context-reconstructed riddle dataset to ensure they preserve the underlying reasoning pattern. See Figure 10 for a 2-shot context-reconstructed prompt in Bengali.
- **Self-evaluation prompt.** We used a single-token self-evaluation prompt to have the model judge candidate answers as correct or incorrect. See Figure 11 for the exact instruction layout.

Prompt : 2-shot Reconstructed Context

You are an expert in solving Bengali riddles.
 Below are some examples. Each example consists of an original Bengali riddle with its answer, followed by one or more 'context-reconstructed' Bengali riddles.
 These reconstructed riddles share the same underlying reasoning pattern as the original but are set in a different context, and their answers are also provided.
 Study these examples carefully to understand how the reasoning pattern is maintained across different contexts.
 For the final riddle presented, provide ONLY the answer. Do not include any explanations, introductions, or any text other than the answer itself.

Examples:
 Example 1:
 Riddle: বুদ্ধি খাটিয়ে বনো দেখি, কোন ফলে ফল ধরে না?
 Answer: বিফল

Example 2:
 Riddle: কোন বিচারে বিচার হয় না?
 Answer: অবিচার

Now, solve this riddle:
 Riddle : যদি থাকে বুদ্ধি বনো দেখি কোন সুখে সুখ নেই?
 Answer :

Figure 10: 2-shot context reconstructed prompt (Bengali). Exemplars are selected from that context-reconstructed riddles of each riddle, drawn from the Bengali riddle dataset.

A.3. Detailed Results

The following tables report the numerical results corresponding to the visualizations shown in Section 4. Tables 3–6 report the values used to compute the aggregated averages shown in the figures.

Strategy	G. Pro	G. Flash	M. Saba	L. Scout	L. Mav	Avg.
0-shot Std.	41.52	27.23	15.25	9.53	13.84	21.47
2-shot Rand.	40.48	27.83	21.13	8.48	15.77	22.74
4-shot Rand.	42.11	28.72	19.79	9.97	15.63	23.24
2-shot Sem.	41.52	28.27	17.86	11.01	15.33	22.80
4-shot Sem.	36.61	26.78	18.45	11.31	15.03	21.64
2-shot Ctx.	39.43	26.49	17.71	10.42	15.33	21.87
4-shot Ctx.	38.39	26.49	19.64	10.57	16.97	22.41

Table 3: Average accuracy (%) across seven Indian languages by prompting strategy and model.

Prompt : Self-Evaluation

You are an expert *{Language}* riddle analyst.
Your task is to evaluate if the provided "Candidate Answer" is a good and correct solution for the given "*{Language}* Riddle".
Critically analyze the riddle and the provided answer to determine if the answer is a correct and fitting solution.
Your response MUST be a single word: either "Yes" or "No".
Do NOT provide any explanations, reasoning, introductory text, punctuation, or any other words. Your entire output must be only "Yes" or only "No".
- Respond "Yes" if the candidate answer is a correct or a very plausible solution.
- Respond "No" if the candidate answer is incorrect, nonsensical, or a poor fit for the riddle.

Riddle: *{Language Riddle}*
Candidate Answer: *{Generated candidate answer}*

Your single-word (Yes or No):

Figure 11: Prompt template for self-evaluation mechanism. The model is instructed to output a single 'Yes' or 'No' token.

Strategy	G. Pro	G. Flash	M. Saba	L. Scout	L. Mav	Avg.
0-shot Std.	88.70	84.84	80.82	79.29	80.24	82.78
2-shot Rand.	87.99	84.79	82.64	78.63	80.76	82.96
4-shot Rand.	88.44	85.07	82.46	79.20	80.39	83.11
2-shot Sem.	86.69	84.73	81.81	80.13	80.59	82.79
4-shot Sem.	85.34	84.50	82.26	79.97	80.29	82.47
2-shot Ctx.	86.14	84.23	81.83	79.46	80.43	82.42
4-shot Ctx.	85.46	84.35	82.46	80.00	81.07	82.67

Table 4: Average BERTScore F1 across seven Indian languages by prompting strategy and model.

A.4. Limitations

Below we discuss some limitations of using riddles for evaluation.

- **Cultural coverage.** Even with seven languages, the benchmark does not capture the full cultural or dialectal variation within each language community. Performance may differ for regional riddles not represented in our dataset.
- **Annotator bias.** Normalizing answers and checking model outputs requires human judgment. This can introduce some subjectivity when deciding whether different synonyms or phrasing should be treated as correct.
- **Scalability.** The dataset construction process relies exclusively on manual curation

Language	G. Pro	G. Flash	M. Saba	L. Scout	L. Mav	Avg.
Bengali	41.37	26.19	17.19	8.78	13.09	21.32
Gujarati	57.14	43.30	29.61	22.92	27.08	36.01
Hindi	44.94	37.50	39.14	16.82	24.55	32.59
Kannada	23.51	17.11	10.27	5.06	8.93	12.98
Malayalam	40.63	21.13	3.42	4.91	9.08	15.83
Tamil	50.00	33.78	22.17	9.08	19.05	26.82
Telugu	22.47	12.80	8.03	3.72	6.10	10.62

Table 5: Average accuracy (%) across seven prompting strategies by language and model.

Language	G. Pro	G. Flash	M. Saba	L. Scout	L. Mav	Avg.
Bengali	87.28	84.24	82.73	80.54	81.04	83.17
Gujarati	91.11	88.80	85.06	83.13	83.22	86.26
Hindi	88.72	87.86	88.52	82.68	83.65	86.29
Kannada	81.98	81.72	79.07	77.56	77.83	79.63
Malayalam	86.69	82.13	76.41	76.49	77.84	79.91
Tamil	89.48	86.11	83.14	79.19	82.20	84.03
Telugu	83.51	81.66	79.35	77.07	77.99	79.92

Table 6: Average BERTScore F1 (%) across seven prompting strategies by language and model.

and review to ensure linguistic correctness and cultural relevance. While this approach improves dataset quality, it limits scalability when extending the benchmark to additional languages. Future work will attempt to explore a semi-automated pipeline with LLMs-in-the-loop, where models assist in generating candidate riddles while human reviewers focus on validation.

- **Training data overlap.** Since many traditional riddles are publicly available online and may have been included in the training corpora of LLMs, there is a possibility of data contamination. Models may therefore have prior exposure to some riddles, which could influence performance.