

Translation, Identity, and Cultural Representation in Indian Literature

Dr. Geeta Rawat

Department of Social Sciences

Shri Guru Ram Rai University

Dehradun, Uttarakhand, Pin Code-248001,

Orcid id- <https://orcid.org/0009-0007-7163-9847>



<http://www.ejset.org/> || Vol. 2 No. 3 (2026): September Issue

Date of Submission: 02-08-2026

Date of Acceptance: 17-08-2026

Date of Publication: 02-09-2026

Abstract— The rapid development of multilingual artificial intelligence has expanded the possibilities for translating Indian literary texts across linguistic and cultural boundaries. However, translation quality in computational research is still assessed primarily through lexical, semantic, or sentence-level correspondence, leaving cultural identity preservation comparatively under-examined. This study proposes a culturally grounded framework for evaluating AI-assisted literary translation from selected Indian languages into English. It focuses on whether translation systems retain identity-bearing elements such as kinship relations, caste and community references, honorifics, culturally embedded metaphors, ecological imagery, idiomatic expressions, and region-specific narrative voice. The proposed research combines multilingual neural and large-language-model translation with computational similarity measures and structured evaluation by bilingual human readers. A Cultural Representation Preservation Index is introduced to distinguish technically fluent translations from translations that maintain culturally meaningful information. The study further examines whether discourse-aware contextual prompting and cultural gloss

information improve representation without producing excessive explanatory translation or semantic distortion.

Keywords— Literary Translation, Cultural Representation, Indian Languages, Large Language Models, Machine Translation, Cultural Identity

Introduction

Translation has played a central role in the circulation of Indian literature because India's literary traditions developed within an exceptionally multilingual cultural environment. Literary expression in Hindi, Bengali, Tamil, Marathi, Urdu, Malayalam, Kannada, Telugu, Assamese and numerous other languages does more than communicate narrative information. Language frequently encodes social position, geography, religious practice, kinship, memory, community identity and historically situated ways of interpreting the world. Consequently, translating Indian literature involves more than replacing linguistic units in one language with approximately equivalent units in another. The translator must negotiate culturally situated meanings whose significance may depend on relationships between speakers, narrative history, local knowledge and conventions that are not directly represented in the target language.

This problem has acquired renewed significance because translation is increasingly mediated by artificial intelligence. Neural machine translation, multilingual language models and generative large language models can now produce fluent multilingual text with little human intervention. In India, the broader technological environment has also changed substantially. The BHASHINI initiative, launched under India's National Language Translation Mission, has expanded technological infrastructure for machine translation, speech recognition and text-to-speech applications across Indian languages. UNESCO's recent assessment of multilingual education in India identifies language technologies and AI-based translation as increasingly important components of linguistic inclusion and digital accessibility. At the international level, UNESCO's Global Roadmap for Multilingualism in the Digital Era similarly stresses that language technologies must support linguistic diversity rather than deepen existing inequalities between digitally dominant and low-resource languages.

The expansion of such systems creates substantial opportunities for Indian literature. Digitally assisted translation could allow works written in regional languages to reach readers who would otherwise have little access to them. It could support literary archives, comparative literature, multilingual education and the preservation of texts from linguistically underrepresented communities. At the same time, literary translation exposes limitations that are less visible when machine translation is evaluated using news reports, administrative documents or ordinary conversational sentences. A system can generate grammatically correct English while altering the cultural positioning of a character, neutralising an honorific, replacing an unfamiliar metaphor with a conventional English expression, or erasing the distinction between different forms of address.

The distinction between linguistic accuracy and cultural representation is therefore fundamental. Consider a culturally specific kinship expression. A machine-generated translation may convert several relationally distinct Indian-language

terms into the generic English words *uncle*, *aunt*, *brother* or *sister*. At the lexical level the output appears reasonable. Within a literary narrative, however, this compression can remove information about age hierarchy, maternal or paternal lineage, intimacy and social obligation. Similar losses occur when caste titles, regional occupations, ritual vocabulary, food practices, ecological references or speech registers are normalised into globally familiar terminology. Each individual modification may appear minor, yet their cumulative effect can reconstruct the cultural world of the text.

Recent NLP research demonstrates that multilingual performance cannot be equated automatically with socially or culturally neutral representation. Ramesh, Gupta and Singh demonstrated measurable gender bias in Hindi-English machine translation and showed why the grammatical characteristics of Hindi require language-specific approaches to bias evaluation. The finding is important beyond gender because it illustrates that structural characteristics of Indian languages interact with learned representations in ways that generic evaluation procedures may fail to capture.




Fig. 1: Bilingualism in Indian Literature

Source:

<https://kuey.net/index.php/kuey/article/download/8028/6022/15570>

More recent work has extended the problem to discourse. IndicDISCO-MT, introduced in 2026, explicitly identifies the

limitations of sentence-level translation assessment and evaluates discourse phenomena such as pronoun resolution and lexical cohesion across eight Indian languages translated into English. This development is particularly relevant to literature, where meanings commonly develop across sentences rather than within isolated segments. A metaphor introduced in one paragraph may acquire significance several pages later; social relations become understandable through repeated forms of address; and narratorial position depends upon lexical and stylistic consistency across extended passages.

Research on low-resource Indian languages presents an additional concern. The 2025 MMLoSo shared task focused on Bhili, Mundari, Gondi and Santali and documented the difficulties caused by limited parallel resources, linguistic diversity, rich morphology and restricted digital representation. Related work examining Hindi dialects such as Awadhi, Bhojpuri, Braj Bhasha and Magahi also demonstrates that dialectal varieties remain substantially less supported by contemporary speech and translation technology than standardised languages. Such technological asymmetry has direct cultural implications because literary identity is often expressed most strongly through regional and non-standardised linguistic forms.

These developments reveal a methodological problem. Most machine-translation evaluations answer questions such as: *Does the generated sentence convey approximately the same semantic content?* or *How closely does the output resemble a reference translation?* Literary translation requires an additional question: *What version of the represented culture survives the translation?*

Research Gap

Although recent studies have improved evaluation of Indic machine translation through fairness analysis, discourse-aware benchmarks, low-resource language datasets and style-sensitive corpora, there remains limited empirical work that measures **cultural-identity preservation as a distinct dimension of AI-assisted literary translation**. Existing

automatic metrics can reward semantic similarity while remaining insensitive to the deletion, generalisation or reinterpretation of culturally meaningful expressions. Conversely, cultural criticism of translation frequently offers rich qualitative interpretation but rarely provides a reproducible computational framework for comparing multiple AI translation systems.

The present research addresses this intersection. It conceptualises cultural representation not as an abstract judgement of whether a translation "sounds Indian," but as a set of observable textual phenomena that can be annotated, compared and evaluated. These include culturally specific entities, kinship relationships, honorific distinctions, social identities, material culture, religious and ritual terminology, landscape references, metaphors, idioms, code-mixing and narrative register. The principal objective is therefore to develop and experimentally examine a **Cultural Representation Preservation framework** capable of identifying cases in which translation remains semantically plausible but culturally reductive.

A second gap concerns context. Literary passages cannot always be translated adequately one sentence at a time. The study therefore distinguishes isolated-sentence translation from passage-aware translation and culturally contextualised translation. This allows the proposed experimental design to test whether giving AI systems broader narrative context and carefully bounded cultural information improves preservation of culturally meaningful elements without encouraging unnecessary additions.

The central research objective is to determine whether multilingual AI translation systems that receive discourse and cultural context preserve the identity-bearing features of Indian literary passages more effectively than systems translating the same material without such context. A related objective is to examine the extent to which conventional translation metrics correspond with human assessments of cultural fidelity. A substantial divergence between these forms of evaluation would indicate that current notions of

machine-translation quality remain incomplete for literary applications.

Literature Review

2.1 Translation as Cultural Representation

Translation theory has long recognised that literary translation cannot be reduced to formal equivalence. Words belong to cultural systems, and literary meaning frequently arises through historical associations, social conventions and local symbolic structures. Translating a culturally marked expression therefore involves a choice among retention, explanation, approximation, substitution and omission. Each strategy affects how the source culture becomes visible to readers of the target language.

This issue is particularly complex in Indian literature because linguistic variation frequently operates as a narrative resource. Writers may alternate between standard language and dialect, combine Persian, Sanskrit, English or regional vocabulary, or use different forms of address to indicate power and intimacy. The resulting text contains layers of social meaning that may not correspond neatly to English categories. A translation can therefore retain plot and denotative meaning while changing character identity and cultural atmosphere.

Cultural representation also involves asymmetry between source and target languages. English has acquired considerable institutional and digital power in the international circulation of literature. Translation into English can increase visibility, but it can simultaneously encourage linguistic normalisation. Culturally unfamiliar constructions may be reformulated according to target-language expectations of readability. In human translation this is usually a conscious or editorially negotiated strategy. In AI-generated translation, normalisation can emerge implicitly from statistical regularities learned during model training.

2.2 Neural Machine Translation and Indic Languages

The development of transformer architectures substantially improved multilingual machine translation, particularly after large-scale pretraining enabled transfer across related languages. Indic-language translation research has consequently moved from independent bilingual models toward multilingual architectures capable of sharing representations across several languages.

Indic-to-English and English-to-Indic research has nevertheless demonstrated uneven performance across language pairs. Appicharla et al., for example, investigated multilingual neural machine translation with romanised vocabularies during the WAT 2021 shared task and observed substantial differences in performance across Indian languages. Such variation reflects several persistent challenges, including the availability of parallel corpora, script diversity, morphological richness and linguistic distance from English.

Low-resource conditions are especially important from the perspective of cultural representation. Digital corpora do not represent all linguistic communities equally. When a model has encountered extensive standard Hindi or English material but relatively little regional, tribal or dialectal literature, unfamiliar cultural expressions may be translated through representations learned from dominant languages. The result may be fluent output with reduced local specificity.

The MMLoSo 2025 shared task illustrates continuing efforts to address this imbalance by developing resources for Bhili, Mundari, Gondi and Santali. Its parallel datasets and systematic evaluation represent significant progress toward technological inclusion. However, translation quality within such benchmarks continues to be evaluated primarily through measures of linguistic correspondence. This is appropriate for general-purpose MT but does not fully address whether culturally important concepts remain recognisable after translation.

2.3 Bias, Identity and Representational Harm

The literature on responsible NLP provides a complementary perspective. Bias research has demonstrated that computational language systems can reproduce patterns associated with gender, social groups and cultural stereotypes. The Hindi–English study by Ramesh et al. is particularly relevant because it adapts bias evaluation to linguistic characteristics specific to Hindi rather than assuming that methods developed for English can be transferred without modification.




Fig. 2: Cultural Identity

Source: https://www.researchgate.net/figure/Model-of-Cultural-Identity_fig2_315497041

This principle can be extended to cultural representation. A culturally responsible evaluation framework should not assume that general multilingual semantic similarity is sufficient for Indian literary material. Features such as honorificity, kinship hierarchy or caste-coded naming systems may require explicit annotation because their cultural function cannot necessarily be inferred from lexical overlap.

BharatBBQ further demonstrates the need for culturally situated evaluation. The benchmark adapts bias assessment to Indian social contexts across eight languages and thirteen social categories rather than relying solely on benchmarks created primarily around Western sociocultural assumptions.

Although BharatBBQ concerns question answering rather than literary translation, its methodological implication is significant: evaluation datasets must encode the cultural structures relevant to the population and linguistic environment under investigation.

The problem can therefore be understood as one of representational fidelity. A translation system may not generate explicitly stereotypical language, yet it may still cause representational harm through omission or homogenisation. When a specific community practice becomes a generic "festival," a culturally significant garment becomes simply "clothes," or different family relationships collapse into one English kinship label, information about the represented community disappears without producing an obviously incorrect sentence.

2.4 Discourse and Context in Literary Translation

Sentence-level evaluation remains one of the most significant methodological limitations of conventional MT assessment. Literary meaning is highly contextual. Pronouns, recurring metaphors, character relationships and shifts in register often depend on information distributed across a passage.

IndicDISCO-MT directly addresses part of this problem by constructing a discourse-oriented benchmark for Indian-language-to-English translation. The benchmark evaluates phenomena including pronoun resolution and lexical cohesion and therefore moves translation evaluation beyond isolated sentences. Its emergence indicates increasing recognition that strong sentence-level performance does not guarantee coherent translation at discourse level.

For literary translation, discourse awareness is necessary but not sufficient. A system may correctly resolve who a pronoun refers to while still neutralising the socially meaningful form through which one character addresses another. Similarly, lexical cohesion can remain intact while metaphorical or culturally symbolic meaning changes. The present study therefore extends the contextual principle from discourse coherence toward **cultural continuity across discourse**.

2.5 Style, Voice and Cultural Specificity

Style represents another dimension that conventional semantic evaluation does not capture adequately. Recent work on style-aware translation for Indian languages has begun to address this limitation. The STAR-IL corpus, introduced in 2026 for multilingual translation of product reviews, demonstrates that domain and colloquial style can be explicitly incorporated into translation datasets and that targeted fine-tuning can improve style-sensitive translation performance. Although product reviews differ substantially from literature, the study establishes an important methodological point: preserving stylistic properties can be formulated as a measurable machine-translation objective rather than treated exclusively as subjective judgement.

Literary voice requires an even richer interpretation of style. Narrative voice can include sentence rhythm, formality, dialect, culturally situated metaphor, lexical repetition and deliberate shifts between linguistic registers. An AI system optimised principally for fluency may remove precisely those irregularities through which a literary character or narrator becomes culturally distinctive.

Research on dialectal Hindi reinforces this concern. Sood, Rajput and Akhtar's investigation of Awadhi, Bhojpuri, Braj Bhasha and Magahi shows the continuing difficulties faced by language technologies when processing linguistic varieties outside standard Hindi. In literature, however, dialect is not merely noise or variation to be standardised; it may be essential evidence of geographical origin, class, intimacy or cultural belonging.

Proposed Methodology

The proposed study adopts a mixed computational–human evaluation framework to examine whether artificial intelligence systems can translate Indian literary passages into English without substantially weakening culturally embedded identity. The design deliberately separates conventional translation quality from cultural-representation

quality because grammatical fluency and semantic similarity do not necessarily indicate successful literary translation.

The investigation is structured around three translation conditions. In the first condition, each literary sentence is translated independently without access to the surrounding passage. This represents the conventional sentence-level machine-translation setting. In the second condition, the system receives the complete literary passage so that character relations, discourse continuity, recurring expressions and narrative context can influence translation. In the third condition, passage context is supplemented with a limited cultural-information prompt containing explanations of culturally specific expressions identified in the source. The purpose of the third condition is not to tell the model how to translate individual sentences, but to determine whether explicit cultural awareness improves preservation of culturally salient content.

The central research objective is to determine whether contextual and culturally informed translation produces measurable improvements in cultural representation compared with isolated sentence translation.

The study addresses the following research questions:

RQ1: To what extent do AI-based translation systems preserve culturally significant expressions in Indian literary texts?

RQ2: Does passage-level contextual information improve cultural preservation compared with sentence-level translation?

RQ3: Does controlled cultural information further improve translation of identity-bearing linguistic elements?

RQ4: Are conventional automatic translation-quality metrics strongly associated with human assessments of cultural fidelity?

RQ5: Which categories of cultural information are most susceptible to loss, generalisation or inappropriate substitution during AI-assisted translation?

The proposed framework treats the source text as containing identifiable cultural units rather than assuming that the entire passage can be represented by a single translation-quality score. Each unit is labelled according to its cultural function. The principal categories include kinship terminology, honorifics and forms of address, social or community references, religious and ritual vocabulary, material culture, food and clothing terminology, ecological and geographical imagery, idiomatic expressions, culturally embedded metaphors, dialectal forms and code-mixed expressions.

Three bilingual annotators with familiarity with the relevant literary tradition would independently identify such expressions. An expression would be retained in the final annotation only when at least two annotators agree that its interpretation depends materially on culturally situated knowledge or social context.

3.1 Cultural Representation Preservation Index

To complement conventional machine-translation metrics, the study proposes a **Cultural Representation Preservation Index (CRPI)**. Rather than evaluating lexical similarity alone, CRPI assesses whether each annotated cultural unit remains functionally recognisable in the English translation.

Each cultural unit is evaluated on four dimensions:

1. **Semantic retention** – whether the core meaning remains present.
2. **Cultural specificity** – whether a culturally distinct term has been retained rather than unnecessarily generalised.
3. **Relational fidelity** – whether social hierarchy, kinship or interpersonal positioning remains intact.
4. **Narrative compatibility** – whether the translation fits the broader literary context without adding unsupported cultural explanation.

Each dimension can be scored from 0 to 2, where 0 indicates loss or serious distortion, 1 indicates partial preservation and

2 indicates strong preservation. The normalised passage-level score can therefore be expressed as:

$$CRPI = \frac{\sum_{i=1}^n (S_i + C_i + R_i + N_i)}{8n}$$

where (n) represents the number of annotated cultural units, (S_i) semantic retention, (C_i) cultural specificity, (R_i) relational fidelity and (N_i) narrative compatibility.

CRPI values consequently range from 0 to 1. Higher values indicate stronger preservation of culturally meaningful information.

An additional **Cultural Generalisation Rate (CGR)** is proposed to identify cases in which culturally specific source expressions are replaced by broader target-language categories. For example, converting several culturally differentiated kinship expressions into the single English term *uncle* would be counted as generalisation even when the translation remains broadly understandable.

The study also records **Cultural Hallucination Incidence (CHI)**, referring to translated information that introduces a cultural explanation, social category, ritual interpretation or contextual claim that is absent from the source passage.

These measures allow culturally responsible translation to be distinguished from both excessive domestication and excessive explanatory expansion.

Experimental Setup

4.1 Dataset Construction

A multilingual literary evaluation corpus is proposed using copyright-permissible, public-domain or appropriately licensed Indian literary texts. The corpus should contain approximately 1,200 passages drawn from six linguistically diverse Indian languages:

- Hindi

- Bengali
- Marathi
- Tamil
- Malayalam
- Urdu

Approximately 200 passages would be selected from each language. Each passage should contain between 80 and 180 source-language words, providing sufficient contextual material while remaining manageable for expert annotation.

Selection would prioritise narrative prose rather than technical, journalistic or administrative writing. Passages would be sampled specifically because they contain at least one culturally meaningful linguistic feature. Text containing only universally interpretable descriptive language would therefore not dominate the benchmark.

To avoid author-specific memorisation effects, passages should be distributed across multiple writers, periods and literary genres wherever licensing permits. Public-domain material should preferably include fiction, memoir, short stories and culturally situated narrative prose.

The corpus would not treat English reference translations as absolute ground truth. Published translations, where legally usable, may serve as one reference layer, but bilingual evaluators would independently assess the relationship between the original passage and generated translations.

4.2 Data Preprocessing

Source texts would first undergo Unicode normalisation and removal of digitisation artefacts. Punctuation would be preserved because punctuation contributes to narrative rhythm and discourse structure.

Language-specific tokenisation would be used only for computational evaluation; the original surface form would remain intact during generation. Transliteration would not

replace native script because transliteration can obscure lexical and morphological distinctions.

Named entities would not be automatically converted into English equivalents during preprocessing. This is important because names, titles and place references may themselves carry cultural information.

The corpus would then undergo manual cultural-unit annotation. Annotators would identify the source span, cultural category, contextual meaning and potential translation risk. Annotators would also indicate whether transliteration, semantic translation, descriptive translation or retention with contextual inference represents an acceptable translation strategy.

4.3 Model Architecture

The experimental comparison would include three broad classes of contemporary multilingual translation technology.

The first category would consist of a dedicated multilingual neural machine-translation model. Such systems provide a useful baseline because they are optimised primarily for translation rather than general-purpose generation.

The second category would consist of a multilingual encoder-decoder architecture capable of transferring knowledge across Indian language pairs.

The third category would consist of instruction-following large language models with demonstrated multilingual capabilities. Unlike conventional MT systems, these models can receive contextual instructions explaining the literary task and cultural constraints.

Rather than treating one specific commercial model as universally representative, the framework is designed to permit substitution as multilingual systems evolve. This makes the proposed methodology usable beyond a single software version.

For large language models, three prompting configurations would be evaluated:

Configuration A: Direct Translation

The model receives only the sentence to be translated.

Configuration B: Passage-Aware Translation

The complete passage is supplied, and the model is instructed to translate the target sentence while maintaining character relationships, register and preceding narrative context.

Configuration C: Cultural-Context Translation

The model receives the passage together with brief source-side annotations explaining potentially ambiguous cultural expressions. The instruction explicitly prohibits adding information that is not supported by the source.

This progression enables the experiment to isolate the effect of discourse context and explicit cultural knowledge.

4.4 Training Strategy

The proposed benchmark is principally an evaluation study rather than a large-scale model-training project. Therefore, the primary experiment would use pretrained multilingual systems without modifying their parameters.

A secondary adaptation experiment may be undertaken using parameter-efficient fine-tuning on culturally annotated parallel passages. Low-Rank Adaptation or another parameter-efficient technique would be appropriate because it allows model behaviour to be adjusted without full retraining.

The adaptation set would be separated at author and text level from the evaluation set. This precaution reduces the possibility that stylistically similar passages from the same source appear in both training and evaluation partitions.

Where model adaptation is conducted, the corpus would be divided into training, validation and testing subsets at approximately 70%, 15% and 15%, with stratification by language and cultural category.

Early stopping would be based on validation performance using a combined semantic and cultural-preservation objective rather than lexical similarity alone.

4.5 Evaluation Metrics

Conventional evaluation would include an established semantic machine-translation metric such as COMET where technically feasible, supplemented by sentence-level semantic similarity.

The proposed cultural evaluation would include:

Cultural Representation Preservation Index (CRPI): preservation of culturally important meaning and social positioning.

Cultural Generalisation Rate (CGR): proportion of culturally specific expressions translated into unnecessarily broad categories.

Cultural Hallucination Incidence (CHI): frequency of unsupported cultural additions.

Narrative Consistency Score (NCS): consistency of character relationships, pronouns, naming and repeated culturally meaningful terminology across the passage.

Register Preservation Score (RPS): human evaluation of whether relative formality, social hierarchy and stylistic voice remain recognisable.

Expert Preference Rate: proportion of passages in which bilingual evaluators prefer one translation condition over another for literary publication.

Inter-annotator reliability should be measured before aggregated cultural scores are reported. Krippendorff's alpha or an appropriate weighted agreement statistic can be used depending on the final rating scale.

Statistical analysis would compare translation conditions using repeated-measures procedures because the same source passages are translated under each condition. When score distributions violate parametric assumptions, Friedman tests

followed by corrected pairwise Wilcoxon signed-rank tests would provide a suitable alternative.

Spearman correlation would be used to examine the relationship between conventional semantic metrics and human cultural-preservation scores.

Results and Discussion

Because the present manuscript proposes a reproducible empirical framework rather than claiming that the described multilingual experiment has already been executed, numerical performance values are intentionally not invented. The results framework below specifies the observations that should be reported after the benchmark is run. This distinction is essential for avoiding unsupported empirical claims.

The principal analysis should first determine whether semantic translation quality and cultural representation behave as equivalent outcomes. The theoretical expectation is that they will overlap only partially. A translation can communicate the principal event described in a passage while reducing culturally relevant social information.

Kinship terminology provides a particularly clear test. Indian languages frequently encode family relationships more precisely than English. When a translation system converts multiple source-language relationships into broad English categories, semantic comprehension may remain relatively high while relational specificity declines. Such cases should receive acceptable semantic scores but lower CRPI and higher CGR values.

Honorifics constitute a second important category. Indian literary dialogue may use pronoun selection, titles, suffixes and lexical forms to establish hierarchy or intimacy. English translations often require indirect strategies for retaining these distinctions. A literal sentence-level system may therefore remove relational information that a context-aware system could recover through vocabulary, tone or appropriate retention of culturally meaningful address terms.

Culturally embedded metaphors are expected to produce a different type of difficulty. Direct lexical translation may preserve the words but not the cultural association. Conversely, aggressive localisation may replace the source metaphor with a target-culture metaphor and thereby improve immediate readability at the expense of cultural specificity. CRPI is designed to penalise both substantial meaning loss and unnecessary cultural replacement.

Passage-aware translation should be especially relevant where the meaning of a cultural unit depends on previously introduced characters or events. For instance, the appropriate interpretation of an address term may not be evident from a single sentence but may become clear from earlier kinship information. Contextual translation can therefore improve relational fidelity even if lexical accuracy remains unchanged.


Culturally contextualised prompting introduces a further trade-off. Explanatory information may help a model distinguish culturally specific concepts, but excessive cultural prompting can encourage over-explication. Literary translation is not an anthropological glossary. Adding explanations such as ritual history or community background directly into narrative sentences may constitute hallucination or stylistic interference when those explanations are absent from the source.

The best-performing cultural translation condition should consequently not be defined as the condition producing the longest or most explanatory text. Strong performance requires preservation with restraint.

Table 1. Proposed Result Interpretation Framework for AI-Assisted Literary Translation

Evaluation Parameter	Sentence-Level Translation	Passage-Aware Translation	Cultural-Context Translation
----------------------	----------------------------	---------------------------	------------------------------

Semantic meaning retention	Establish baseline	Examine contextual improvement	Examine whether cultural information affects core meaning
Cultural Representation Preservation	Identify culturally lost information	Test recovery through discourse context	Test preservation with explicit cultural support
Cultural Generalisation	Detect replacement of specific concepts with broad English categories	Determine whether context reduces generalisation	Assess whether annotations preserve distinctions
Cultural hallucination	Record unsupported additions	Examine contextual inference errors	Detect over-explanation caused by cultural prompting
Narrative consistency	Identify sentence-level inconsistencies	Evaluate continuity across passages	Determine whether cultural context strengthens continuity
Register preservation	Assess loss of hierarchy and voice	Examine context-sensitive tone	Examine whether cultural guidance retains



The proposed table deliberately avoids reporting fictitious numerical measurements. Once the experiment is completed, the qualitative entries should be replaced with observed values, confidence intervals and statistical significance where appropriate.

5.1 Relationship Between Automatic and Cultural Metrics

One of the most significant expected contributions concerns metric disagreement. Conventional automatic evaluation systems generally focus on semantic correspondence, reference similarity or learned judgements of translation quality. These measurements are useful but cannot be assumed to capture all culturally salient transformations.

A culturally specific noun translated into a generic English term may remain semantically close enough to receive a strong automatic score. A human evaluator familiar with the source culture may nevertheless identify substantial loss.

The correlation analysis between semantic metrics and CRPI is therefore crucial. A very high correlation would suggest that existing translation-quality measures already capture much of the culturally relevant information. A moderate or weak relationship would provide empirical justification for cultural evaluation as an independent dimension.

5.2 Translation of Identity-Bearing Expressions

The most analytically valuable outputs are likely to involve expressions whose function depends on identity or social relation rather than dictionary meaning.

Caste, community, occupational and regional references require particular care. The objective should not be to amplify social categories that the author does not emphasise, but neither should translation erase categories that structure the narrative.

Similarly, religious vocabulary should not automatically be substituted with concepts from other traditions merely because a superficially similar target-language term exists. Retention or transliteration may sometimes provide greater representational accuracy.

The same principle applies to food, clothing and material culture. A culturally specific item may be unfamiliar to English-language readers, yet unfamiliarity alone does not justify replacing it with a globally familiar object. Literature often gains meaning precisely through culturally located material details.

5.3 Implications for Responsible Artificial Intelligence

The study extends the concept of responsible AI beyond conventional concerns such as toxic generation, explicit stereotypes and demographic bias. Cultural homogenisation represents another form of model behaviour deserving systematic examination.

An AI system can be polite, fluent and factually plausible while repeatedly translating minority or regional cultural distinctions into dominant-language categories. At scale, such behaviour could influence how digitally translated literature represents communities to global readers.

This issue becomes increasingly important as automated translation enters publishing, digital libraries and educational systems. Human literary translators traditionally make interpretative choices and may justify those choices through editorial practice. Automated systems can perform similar

transformations without exposing the assumptions that produced them.

Cultural evaluation therefore has implications for model transparency. Publishers and translators using AI-assisted workflows should know not only whether a model generates fluent English, but also which types of cultural content it tends to generalise.

Future Scope

Future research can extend the proposed framework from six languages to low-resource and tribal languages whose literary traditions remain underrepresented in computational datasets. Bhili, Gondi, Santali, Mundari and numerous regional varieties represent particularly important directions because cultural preservation and digital linguistic inclusion are closely connected.

A second extension could investigate translation between Indian languages rather than only translation into English. Hindi-to-Tamil, Bengali-to-Marathi or Malayalam-to-Hindi literary translation may reveal different cultural-transfer patterns because both source and target languages belong to culturally heterogeneous Indian environments.

Future systems could also incorporate retrieval-augmented translation. Instead of relying exclusively on information encoded in model parameters, a translation model could retrieve an author-specific glossary, historical context, character relationships and previously translated terminology. Such retrieval could improve consistency without requiring the model to invent cultural explanations.

Cultural evaluation could further be extended to poetry, drama, oral narratives and folk literature. These genres pose additional problems involving meter, performance, sound symbolism and non-standard language.

Another promising direction is culturally informed human–AI collaborative translation. Rather than replacing the translator, an AI system could identify potentially sensitive cultural units and present alternative translations together

with warnings when a culturally specific expression has been generalised.

Finally, CRPI could eventually be developed into an evaluation benchmark covering multiple cultures and language families. Such work would help multilingual AI research move beyond the assumption that semantic equivalence alone constitutes adequate translation.

Conclusion

Artificial intelligence has considerably increased the technical feasibility of translating literature across languages, but fluent translation does not necessarily imply culturally faithful translation. Indian literary texts present a particularly important test because language frequently carries information concerning family relationships, hierarchy, community, religion, region, material culture and collective memory.

This study proposes a framework for evaluating these dimensions systematically. Its central contribution is the distinction between semantic translation quality and cultural-representation preservation. The proposed Cultural Representation Preservation Index evaluates whether culturally meaningful units retain their semantic content, specificity, relational function and narrative compatibility after translation.

The experimental framework further distinguishes sentence-level translation, passage-aware translation and culturally contextualised translation. This structure allows researchers to determine whether discourse information and carefully bounded cultural knowledge improve translation without encouraging unsupported explanation.

The study also introduces complementary measures of cultural generalisation, hallucination, narrative consistency and register preservation. Together, these measures provide a more detailed understanding of literary translation than conventional similarity metrics alone.

The proposed approach does not argue that artificial intelligence should determine a single correct representation of Indian culture. Literary translation remains interpretative, and culturally meaningful expressions may support multiple defensible renderings. The purpose of computational evaluation is instead to make cultural loss more visible and measurable.

As multilingual AI systems become increasingly involved in publishing, education, archives and cross-cultural communication, evaluation frameworks must address not only what information survives translation but also whose language, identity and cultural specificity remain visible. For Indian literature, responsible AI-assisted translation should therefore aim for more than grammatical English or semantic approximation. It should preserve enough of the source text's cultural structure for readers to encounter the literary world on its own terms rather than through unnecessary linguistic and cultural homogenisation.